

**IN THE UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF TEXAS
WACO DIVISION**

SONOS, INC.,
Plaintiff

No. 6:20-cv-881-ADA

v.

GOOGLE LLC,
Defendant

**DR. KYRIAKAKIS' DECLARATION IN SUPPORT OF GOOGLE LLC'S RESPONSIVE
CLAIM CONSTRUCTION BRIEF**

I. INTRODUCTION

1. My name is Christos Kyriakakis. I am currently an Associate Professor of Electrical and Computer Engineering at the University of South California. I have been retained by Google, Inc. ("Google") to provide opinions concerning the claim construction and indefiniteness of certain limitations in the asserted claims of U.S. Patent Nos. 9,344,206; 10,469,966; and 10,848,885 (the "Zone Scene" Patents)¹ and U.S. Patent Nos. 9,967,615 and 10,779,033 Patents (the "Direct Play" Patents),² (collectively, the "Asserted Patents").

2. For my work on this case, I am being compensated for my time at my typical consulting rate of \$525 per hour. I am also being reimbursed for all incurred expenses. My compensation does not depend on the substance of my opinions or the outcome of any issues in this case. I have no other interests in this litigation or with any of the parties.

¹ These three patents are each continuations, sharing substantively the same specification and all claim priority to a single provisional application filed on Sept. 12, 2006.

² These two patents are continuations, both claim priority to an application filed on Dec. 30, 2011.

3. My opinions are set forth below. I understand that Sonos may submit arguments in support of its proposed constructions and I reserve the right to supplement my opinions in response to those arguments.

4. I have personal knowledge of the facts and opinions set forth in this declaration, and, if called upon to do so, I would testify competently thereto.

5. In rendering my opinions, I have considered the Asserted Patents, their file histories, and any other documents referenced, discussed, or listed in my declaration, and my own knowledge and experience in the fields such as wireless communications, and multimedia systems. I have also reviewed the productions associated with the claim construction disclosures in this case, including the extrinsic and intrinsic evidence cited by the parties. A list of the materials I considered is set forth in **Exhibit B**.

6. In forming my opinions, I understand that the claims should be interpreted as they would be understood by a person of ordinary skill in the art of the patents at the time its application was filed. I understand that the claims are to be construed with reference to the patent's specification, the claims, the prosecution history, in light of the plain meaning of the terms used in the claims, and with reference to other sources of information, such as dictionaries, textbooks, and literature or other patents in the same or related fields.

7. My analysis of the materials produced in this matter is ongoing and I will continue to review any new material as it is provided. This declaration represents only those opinions I have formed to date. I reserve the right to amend or supplement my opinions based on additional documents or evidence I am presented, including without limitation any arguments or expert declarations advanced by Sonos in this case.

II. BACKGROUND & QUALIFICATIONS

8. In this section I have summarized my education, career history, publications, and other relevant information. My curriculum vitae, which includes my qualifications as well as my publications, is attached as Exhibit A.

A. Educational Background

9. I earned my Bachelor of Science degree in Engineering and Applied Science from the California Institute of Technology (Caltech) in 1985. I received my Master of Science degree in Electrical Engineering in 1987 and my Ph.D. in Electrical Engineering in 1993, both from USC. My expertise is in audio and acoustic sciences. My research interests lie at the intersection of acoustics, psychoacoustics (the science that studies human perception of sound), and audio signal processing. My recent research has focused on the study of audio systems in challenging environments including automobiles and mobile devices, as well as algorithms for enhancing the performance of voice recognition engines. I have published several technical papers on acoustical measurement and calibration methods that can be applied to listening rooms, movie theaters, headphones, and automobiles, and developed novel signal processing algorithms for optimizing sound system performance of speakers and headphones. Other topics I have researched include multichannel audio streaming over high bandwidth networks, audio acquisition and rendering, virtual microphones and virtual speakers, hybrid headphone-loudspeaker rendering methods, and advanced signal processing techniques for optimizing sound quality in automobiles.

B. Relevant Professional Experience

10. I am the founding Director of the USC Immersive Audio Laboratory with facilities for experimental work in room acoustics, multichannel audio, and psychoacoustics.

This laboratory also serves as a unique teaching facility for my undergraduate course in Introduction to Digital Audio and my graduate course in Immersive Audio Signal Processing. Both courses have a major acoustics component that examines the interaction of sound with the acoustical environment (home, movie theater, car). The graduate course was developed through a two-year grant I received from the National Science Foundation entitled “Collaborative Learning in Engineering Using Immersive Environments,” and was the first of its kind to assess the impact of audio immersion in student learning. In addition to the courses I teach, I have also supervised and served on Ph.D. dissertation committees for more than 30 students.

11. From 2003-2018, I was also the founder and Chief Technology Officer of Audyssey Laboratories, a USC spin-off company that develops and licenses audio technology to leading automotive, professional and consumer electronic companies around the world including Jaguar, Land Rover, Audi, Mercedes Benz, Volvo, IMAX, Denon and Intel. As part of my work at Audyssey, I developed acoustical measurement methods for characterizing the performance of headphones and earbuds. Using the system I designed, I led the effort to create a database of 500 headphone measurements from nearly 100 manufacturers. I also led the development of audio algorithms and designed speakers with Bluetooth and WiFi playback capability. These speakers were novel acoustical designs that used a combination of unique enclosures, and audio signal processing to optimize their performance and overcome limitations that arise from small drivers and enclosures. For example, we used signal processing technologies combined with novel acoustical design to extend the bass response of small woofers and passive radiators beyond what was previously possible in small speaker enclosures. The innovations in these designs have received awards, including Popular Science’s “Best of What’s New.”

12. In April 2018 I was elected as Senior Member of the Institute of Electrical and Electronic Engineers in recognition of my contributions to the field of engineering. I am also a member of the Audio Engineering Society, the leading association for professionals in the audio industry. I have published nearly 100 peer reviewed technical papers. I have published a book entitled Immersive Audio Signal Processing, and hold several patents in acoustic measurement of loudspeakers in rooms and cars, headphone optimization, loudspeaker crossover optimization, and loudspeaker response correction using signal processing. My publications examine various aspects of sound measurement, how sound interacts with the acoustical elements of the environment, novel methods for surround sound recording and reproduction, and the perception of sound by human listeners. In 2006, I received a World Technology Network Award. This organization presents awards to innovators in several areas in which technology can foster a paradigm change. My award was for innovations in immersive audio that enable new capabilities in media and journalism. Other award recipients at that event included Vice President Al Gore, Google, and Space-X.

13. In the late 1990s and early 2000s, I was a faculty researcher and later Deputy Director of the National Science Foundation's engineering research center established at USC. I was studying the fundamental and technological limitations of immersive audio and the role of acoustics on the performance of loudspeakers and audio systems in homes and cars. In 2003, together with one of my graduate students, I received the award for Best Paper at the Institute of Electrical and Electronics Engineers ("IEEE") Conference on Signals, Systems and Computers. In my role as Deputy Director, I lead a team of faculty and student researchers in developing the first distributed immersive performance system that included streaming architectures for synchronizing multiple high definition and multichannel audio streams.

III. LEGAL STANDARD

14. I am not a lawyer, and I do not intend to offer any opinions as to the interpretation of the law. In this section, I describe my understanding of certain legal standards. I have been informed of these legal standards by Google's attorneys. I am relying only on instructions from Google's attorneys for these legal standards. I set forth my understanding below.

A. Person of Ordinary Skill in the Art ("POSITA")

15. I understand that claim construction and indefiniteness are both analyzed from the perspective of a person having ordinary skill in the art. I understand that the person of ordinary skill in the art is a hypothetical person or ordinary creativity, not an automaton. I understand that a person of ordinary skill, while not someone who undertakes to innovate, is capable of drawing inferences and taking creative steps. I understand that, in determining the level of skill in the art, courts consider the type of problems encountered in the art, prior art solutions to those problems, rapidity with which innovations are made, sophistication of the technology, and the educational level of active workers in the field. I understand that not all of these factors will be relevant in a given case.

B. Patent Claims and Claim Construction

16. I understand that a patent may include two types of claims, independent claims and dependent claims, that an independent claim stands alone and includes only the limitations it recites, that a dependent claim can depend from an independent claim or another dependent claim. I understand that a dependent claim includes all the limitations that it recites in addition to all of the limitations recited in the claim from which it depends.

17. I understand that the words of a claim are generally given their ordinary and customary meaning. I understand the ordinary and customary meaning of a claim term is the

meaning that the term would have to a person of ordinary skill in the art in question at the time of the invention. I understand that the person of ordinary skill in the art is deemed to read the claim term not only in the context of the particular claim in which the disputed term appears, but in the context of the entire patent, including the specification.

18. I understand there are only two exceptions to the general rule that words of a claim are given their plain and ordinary meaning: first, when a patentee sets out a definition and acts as his own lexicographer; and second, when the patentee disavows the full scope of a claim term either in the specification or during prosecution. I understand that these standards are exacting. I understand that, to act as its own lexicographer, a patentee must clearly set forth a definition of the disputed claim term other than its plain and ordinary meaning and clearly express an intent to redefine the term. I understand that disavowal requires a clear and unmistakable disclaimer. I understand that, absent disavowal or lexicography, it is improper to import limitations into the claims from the patent specification, or to limit the claims to a particular embodiment. I understand that courts may consider extrinsic evidence outside of the patent and its file history, such as dictionaries, scientific treatises, and testimony from experts and inventors. However, I also understand that extrinsic evidence is less significant than the intrinsic record in determining the legally operative meaning of claim language.

C. Indefiniteness

19. I understand that 35 U.S.C. § 112, ¶ 2 requires patent claims to be definite, meaning a claim should “particularly point[] out and distinctly claim[] the subject matter that the applicant regards as the invention.” I understand that indefiniteness is to be evaluated from the perspective of someone skilled in the relevant art at the time the patent was filed. I understand that, in assessing indefiniteness, claims are to be read in light of the patent’s specification and

prosecution history. I understand that the definiteness requirement of § 112, ¶ 2 requires that a patent's claims, viewed in light of the specification and prosecution history, inform those skilled in the art about the scope of the invention with reasonable certainty. I understand that the claims, when read in light of the specification and the prosecution history, must provide objective boundaries for those of skill in the art.

20. I am also informed that the specification of a patent must satisfy a definiteness requirement, which requires that it conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as the invention.

21. I am also informed that definiteness requires that a patent's claims, viewed in light of the specification and file history from the perspective of a person skilled in the relevant art at the time the patent was filed, inform those of ordinary skill in the art about the scope of the invention with reasonable certainty.

22. I understand that a patent must be precise enough to afford clear notice of what is claimed and apprise the public of what is still open to them in a manner that avoids a zone of uncertainty which enterprise and experimentation may enter only at the risk of infringement claims.

IV. THE SONOS PATENTS

A. The Direct Play Patents


23. The "Direct Play" patents claim priority through a series of continuation applications to U.S. Application No. 13/341,237, filed on December 30, 2011, and name Tad Coburn and Joni Hoadley as inventors. The common specification of the Direct Play patents discloses allowing a user to switch playback of remote content from a computing device (e.g., laptop or mobile phone) to a playback device (e.g., speaker, multimedia unit such as a television,

etc.) in a local network. '033 Pat., 1:20-40, 2:8-29, 3:28-30, 12:44-67, 13:14-22, 13:54-56. The computing device or controller (terms that are used interchangeably in the patents) provide a graphical user interface (GUI) for “navigat[ing] a playlist of multimedia items and control[ing] operations of one or more [playback devices].” '033 Pat., 3:30-33, 9:10-48, 11:43-45. A local playback network includes playback devices and a computing device that may be connected to a non-local media source. *Id.*, 11:62-12:67. This allows either device to fetch content from a third-party service. *Id.* When a playback device is configured for playback of content, it can retrieve media by using a uniform resource indicator (URI) “that specifies an address to particular [content] in the cloud.” *Id.*, 11:62-12:3, 12:53 (noting that a uniform resource locator (URL) is an example of a URI), 13:13-22, 13:54-56, 1:19-29, 8:63-66.

B. The Zone Scene Patents

24. Each of the Zone Scene Patents originated with U.S. Provisional Application No. 60/825,407 (“the ‘407 Provisional”), which was filed on September 12, 2006. The ‘966 and ‘855 patents are both direct continuations of U.S. Non-Provisional Application No. 15/130,919 (the “‘919 Application”), which was filed on April 15, 2016, and the ‘919 Application is a direct continuation of the ‘206 patent. Each of the Zone Scene patents name Robert Lambourne as an inventor. For convenience, I refer to the column and line numbers of the ‘206 patent’s specification, and substantially the same disclosures appear in the ‘966 and ‘885 patents.

25. The ‘206 patent discloses a “multi-zone system” with “zone players.” *See, e.g.*, ‘206 Patent at 2:28-37; FIG. 1. Figure 1 of the patent discloses a diagram of an “exemplary embodiment” of the alleged invention:



26. The specification describes Figure 1 as follows:

The configuration may represent, but not be limited to, a part of a residential home, a business building or a complex with multiple zones. There are a number of multimedia players of which three examples 102, 104 and 106 are shown as audio devices. Each of the audio devices may be installed or provided in one particular area or zone and hence referred to as a zone player herein. '206 Pat. at 4:29-36.

27. The Zone Scene patents are directed to a particular type of speaker usage that includes what the patents refer to as a “zone scene”:

If the user wishes to link 5 of the 6 zone players using the current mechanism, he/she must start with a single zone and then manually link each zone to that zone. This mechanism may be sometimes quite time consuming. According to one embodiment, a set of zones can be dynamically linked together using one command. Using what is referred to herein as a theme or a zone scene, zones can be configured in a particular scene (e.g., morning, afternoon, or garden), where a predefined zone grouping and setting of attributes for the grouping are automatically effectuated. '206 Pat. at 8:19-28.

28. The patent describes multi-zone audio systems as conventional in the art, and I agree that many aspects of these speaker systems are conventional:

Currently, one of the systems that can meet part of such demand is a conventional multi-zone audio system that usually includes a number of audio players. Each of the audio players has its own amplifier(s) and a set of speakers and typically installed in one place (e.g., a room). In order to play an audio source at one location, the audio source must be provided locally or from a centralized location. '206 Pat. at 1:40-46.

In my experience, at the time that the Zone Scene patents were filed, multi-zone audio systems existed from a variety of manufacturers, such as Bose, Crestron, and others.

29. The patents describe “zone scenes” as above and in figures, such as Figures 5A and 8 (from the '966 patent) below.

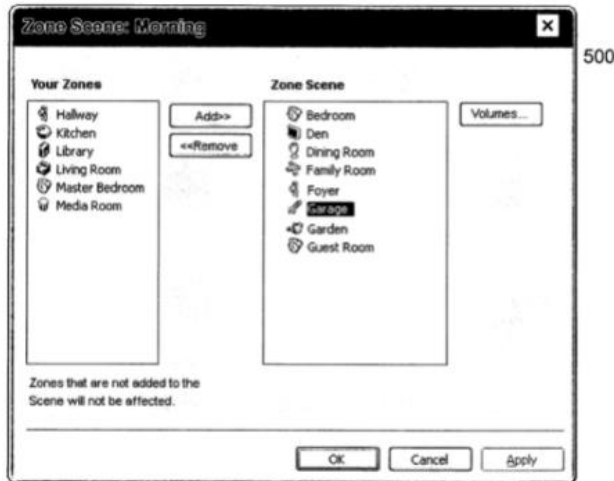


FIG. 5A



FIG. 8

“Zone scenes” are described as providing additional functionality, capabilities, and requirements over “zones” or groups of speakers that were acknowledged as conventional at the time.

E.g., '206 Pat. at 3:42-60; 8:19-9:16; 9:55-10:20.

V. LEVEL OF ORDINARY SKILL IN THE ART

30. I understand that the Zone Scene patents are each continuations, sharing substantively the same specification and all claim priority to a single provisional application filed

on Sept. 12, 2006 and that Sonos contends the Zone Scene Patents are entitled to a priority date of Sept. 12, 2006. I understand that the Direct Play patents are each continuations, sharing substantively the same specification and all claim priority to a single provisional application filed on Dec. 30, 2011 and that Sonos contends the Direct Play Patents are entitled to a priority date of Dec. 30, 2011. For purposes of this declaration only, I have assumed that the asserted claims of the Asserted Patents are entitled to a priority date of Sept. 12, 2006 for the Zone Scene Patents and Dec. 30, 2011 for the Direct Play Patents.

31. In my opinion, a person of ordinary skill in the art at this time would have had a bachelor's of science in electrical engineering, computer science or engineering, or a related field, and two to four years of work or research experience in the field of information networks, data communications, or multimedia systems, or a Master's degree and one to two years of experience in the same field.

32. My determination of the appropriate level of skill is based on my review of the asserted patents, the type of problems encountered in the art, prior art solutions to those problems, the sophistication of the technology, and the educational level of active workers in the field.

33. I have set forth in Section IV above the general subject matter of the asserted patents.

34. I meet the above criteria for a POSITA and consider myself a person with at least ordinary skill in the art pertaining to the asserted patents. I would have been such a person at the time of the alleged invention of each of the Asserted Patents. I am qualified to provide opinions concerning what a POSITA would have known and understood at that time, and my analysis and

conclusions in this declaration are from the perspective of a POSITA at least as of the earliest filing date of the asserted patents.

VI. ANALYSIS

35. I understand that the following claims are asserted in this lawsuit:

Patent	Claims
9,344,206	1-5, 7, 10-19
10,469,966	1-4, 6-12, 14-20
10,848,885	1-3, 5-10, 12-17, and 19-20
9,067,615	1-3, 6-9, 11-15, 18-21, 23-26, 28-29
10,779,033	1-2, 4, 7-13, 15-16

36. I have been asked to provide my opinions as to the proper constructions of certain terms in the asserted claims based on the claim construction principles articulated above. My opinions are set forth below.

VII. CLAIM TERMS

A. “zone configuration” [’206 Patent]

Sonos Construction	Google Construction
Construed “zone configuration characterizes one or more zone scenes” as “configuration data that provides an indication of one or more zone scenes”	“Zone configuration” is indefinite.
No separate construction for “zone configuration”	

B. “group configuration” and “zone configuration” [’206 Patent]


Sonos Construction	Google Construction
No construction necessary, no separate construction necessary	“Group configuration” is indefinite, “Zone configuration” is indefinite

37. I have been asked to opine on the terms “zone configuration” and “group configuration” as recited in the claims of the ’206 Patent.

38. Claim 1 of the ’206 Patent reads, in pertinent part: “wherein the zone configuration characterizes one or more zone scenes, each zone scene identifying a group configuration associated with two or more of the plurality of independent playback devices.” Independent claims 12 and 17 include nearly identical language.

39. In my opinion, the terms “zone configuration” and “group configuration” are indefinite because a POSITA would not understand how to distinguish a “zone configuration” from a “group configuration” or a “zone scene.” In my opinion, neither the ’206 Patent, nor the intrinsic record, nor the general knowledge or skill in the art would provide a POSITA with a sufficient basis for understanding these terms in order to determine whether they would or would not be practicing any of these independent claims. As explained below, the patent refers to “zone configuration” and “group configuration” confusingly, and there is no significant distinction between the terms “zone configuration” and “group configuration” in this field.

40. Beginning with Figure 3A (below), this figure depicts gathering different “zones” on the left into “zone configuration/scenes” on the right. The figure confuses the distinct terms “zone configuration” and “zone scene” by putting a slash between them and not explaining any difference in their meaning.

**FIG. 3A**

41. Further, the patents regularly use the term “zone group,” which further confuses the distinction between “zone” and “group” as well as the distinction between “zone configuration” and “group configuration.”

- Fig. 3A above
- FIG. 3A provides an illustration of one zone scene, where the left column shows the starting zone grouping—all zones are separate, the column on the right shows the effects of grouping the zones to make a group of 3 zones named after “Morning” ’206 Pat. at 3:42-46.
- Two or more zone players may be grouped together to form a new zone group. Any combinations of zone players and an existing zone group may be grouped together. In one instance, a new zone group is formed by adding one zone player to another zone player or an existing zone group. *Id.* at 5:4-8.
- In one embodiment, the controller 240 is used to synchronize more than one zone players by grouping the zone players in a group. In another embodiment, the controller 240 is used to control the volume of each of the zone players in a zone group individually or together. *Id.* at 6:27-31.
- It can be understood that, in a multi-zone system, there may be multiple audio sources being played respectively in more than one zone players. The music transport functions described herein shall apply selectively to one of the sources when a corresponding one of the zone players or zone groups is selected. *Id.* at 7:1-6.
- In one embodiment, an application module is configured to facilitate grouping a number of selected zone players into a zone group and synchronizing the zone players for one audio source. In another

embodiment, an application module is configured to control together the audio volumes of the zone players in a zone group. *Id.* at 7:16-21.

- In one embodiment, a user creates a zone group including at least two zone players from the controller 240 that sends signals or data to one of the zone players. *Id.* at 7:38-40.
- According to one implementation, an application module is loaded in memory 282 for zone group management. *Id.* at 7:51-52.
- The user may start grouping zone players into a zone group by activating a “Link Zones” or “Add Zone” soft button, or de-grouping a zone group by activating an “Unlink Zones” or “Drop Zone” button. The detail of the zone group manipulation will be further discussed below. *Id.* at 7:58-62.
- Using what is referred to herein as a theme or a zone scene, zones can be configured in a particular scene (e.g., morning, afternoon, or garden), where a predefined zone grouping and setting of attributes for the grouping are automatically effectuated. *Id.* at 8:24-28.
- FIG. 3A provides an illustration of one zone scene, where the left column shows the starting zone grouping—all zones are separate, the column on the right shows the effects of grouping the zones to make a group of 3 zones named after “Morning”. *Id.* at 8:32-36.

42. The patents likewise use term the term “zone group configuration.” It is ambiguous whether the term “configuration” is modifying “zone” or modifying “group,” again rendering a POSITA unable to reasonably differentiate between them.

- According to one aspect of the present invention, a software module implementing one embodiment of the present invention is executed, the processor 204 operates in accordance with the software module in reference to a saved zone group configuration characterizing a zone group created by a user, the zone player 200 is caused to retrieve an audio source from another zone player or a device on the network. *Id.* at 5:43-50.
- Typically, a saved zone group configuration file is transmitted to a controller (e.g., the controlling device 140 or 142 of FIG. 1, a computer, a portable device, or a TV) when a user operates the controlling device. The zone group configuration provides an interactive user interface so that various manipulations or control of the zone players may be performed. *Id.* at 5:54-58.
- In another embodiment, a saved zone group configuration is transmitted between a zone player and a controller via the RF interfaces. The

controller 270 may control one or more zone players, such as 102, 104 and 106 of FIG. 1. *Id.* at 7:31-35.

43. The prosecution history of these patents likewise does not provide guidance as to the meaning of the terms “zone configuration” and “group configuration.” These terms are not defined in the prosecution history nor described in any detail therein.

44. The claims also provide no guidance to one of skill as to the differences between “zone configuration” and “group configuration.” For example, claim 1 of the ’206 patent states in pertinent part, “wherein the zone configuration characterizes one or more zone scenes, each zone scene identifying a group configuration associated with two or more of the plurality of independent playback devices.” Accordingly, the zone configuration “characterizes” a zone scene and the zone scene “identifies” a group configuration. It is unclear how the terms “characterize” and “identify” are different as used in these claims. I have reviewed Sonos’s opening claim construction brief and Sonos does not articulate the differences between these terms in this context either. Rather, Sonos generally relies on the fact that it proposed different constructions containing these terms to argue that they are different. That does not answer the question and does not convince me that a POSITA would have a reasonable certainty as to the scope of these claims.

45. Sonos also claims that “the ‘zone configuration’ comprises a data representation of the ‘zone scene.’” Br. at 11. Sonos does not articulate what would be included and excluded from a “data representation” in this context. Sonos also provides no explanation for what, in the computer and speaker field, is *not* a “data representation.” In my view, in the context of speakers and controllers as described in the Sonos patents, there is no substantive difference between a “zone scene” and a “data representation” of a zone scene. Sonos provides no examples or any explanation for this statement. Sonos appears to imply that “zone scene” has

the same information, but in the *abstract*, as “zone configuration.” To the extent that Sonos does argue “zone scene” is some abstract concept, this does not provide reasonable guidance to a person of ordinary skill in the art as to the scope of the claims or how to avoid them because I understand an abstract concept is both patent ineligible and ambiguous in this context.

46. Further, Sonos wrote in its opening brief that a “zone scene” is a “previously saved grouping.” Br. at 2. It is unclear how the “zone scene” could be an abstract concept (as Sonos seems to contend) but also something that was “previously saved.” Sonos does not attempt to explain how something “previously saved” is not a “data representation,” and I have no reasonable certainty as to the meaning of such a term either.

47. As noted in my background and qualifications, I have also designed and worked with speaker systems throughout my career. The terms “zone configuration” and “group configuration” are not well defined in the art and the distinction between those terms is ambiguous in the art as well. These terms were not used to describe speaker systems with which I am familiar in the ways used in the claims, nor is there any discernable difference between those terms in either commercial systems or in the literature in this field.

48. Accordingly, in my opinion the terms “zone configuration” and “group configuration,” viewed in light of the specification and prosecution history, do not inform those skilled in the art about the scope of the invention with reasonable certainty.

C. “local area network” [’615 Patent]

Sonos Construction	Google Construction
“data network that interconnects devices within a limited area, such as a home or office”	Plain and ordinary meaning; no construction necessary at this time

49. It is my opinion that “local area network” requires no construction. Based on my review of the specification for the ’615 patent, the term “local area network” does not have any special meaning.

50. Sonos argues that a LAN is necessarily restricted to transferring “digital data packets,” but I disagree. *See* Br. at 19-20 (and embedded definition of “data network” at 23-26). A LAN is merely a local area network, and like most networks, it can be analog or digital. As I described with respect to the terms data networks and digital data packets below, and which I fully incorporate herein by reference, many networks have been and remain analog and are not limited to transmitting or receiving “digital data packets.”

51. I disagree with Dr. Schmidt’s statements that a POSITA would have understood at the time of the invention that devices such as an infrared remote sending a signal to a TV or devices connected via an Internet connection) would not be coupled by a “local area network.” Br. Ex. 24 at ¶88. Nonetheless, in my opinion a POSITA would have understood the meaning of a “local area network” in the ’615 patent to be no different than in common parlance at the time of the invention.

52. Dr. Schmidt glosses over the intrinsic evidence and spends several paragraphs discussing extrinsic evidence. Br. Ex. 24 at ¶90-91, 96-97. When Dr. Schmidt does mention the ’615 specification, he focused on Figure 1 of the ’615 Patent and related passages but the passages he refers to either discuss a user’s home, or a local area network. Dr. Schmidt has not pointed to any passage in the ’615 specification which equates local area network with a limited area such as a home or office. *See, e.g.*, ’615 Patent at FIG. 1; 5:12-28, 12:44-49, 10:64-66. Br. Ex. 24 at ¶92-94. Similarly, Dr. Schmidt refers to passages of the ’615 specification as allegedly emphasizing the contrast between “local” and “‘cloud,’ ‘remote,’ and ‘Internet,’” but none of the

passages he points to actually contrast a local area network to “‘cloud,’ ‘remote,’ and ‘Internet.’” Br. Ex. 24 at ¶95. The ’615 patent also includes disclosures teaching that a LAN spanning a “hotel” (’615 patent at 13:41-49), which a POSITA would understand as larger than a home or office and typically include multiple buildings spread across a relatively large area. One other asserted patent also discloses that a LAN may span a “complex,” which a POSITA would understand, similar to a hotel, to be larger than a home or office. See ’206 patent at 6:12-16.

53. Accordingly, a POSITA would understand that a LAN can span buildings, office complexes, and larger areas than that of a single home or office. *E.g.*, Ex. C (Deploying the World’s Largest Campus IEEE 802.11b Network) at 4 (showing 802.11 network that was “campus wide,” “150 buildings,” “600 acres coverage”).

D. “a media particular playback system” [’615 Patent]

Sonos Construction	Google Construction
“a media playback system”	Indefinite

54. In my opinion the term “media particular playback system” is indefinite, because it has multiple reasonable interpretations, and a POSITA would not be able to determine which of those reasonable interpretations is correct. The claims and the specification fail to inform a POSITA of the scope of the term “media particular playback system.”

55. I understand that Claims 3, 15 and 26 recite “detecting a set of inputs to transfer playback from the control device to a particular zone group of a media particular playback system.” In contrast, claims 1-2, 13-14 and 25 recite a “particular playback device,” and claims 2 and 14 recite a “media playback system” separate and distinct from the “particular playback device.” In my opinion a POSITA would not readily understand that the inclusion of the word

“particular” in the phrase “media particular playback system” was a typographical error to be corrected by omitting “particular.”

56. I see no indication whether the “media particular playback system” of claims 3, 15, 26 corresponds to a playback system that can only play particular media formats, particular media types, or is a typographical error. A POSITA could theoretically understand “media particular” as intended to differentiate from “multimedia,” *i.e.* a playback system specific to a particular type of media – such as audio.

57. Certain playback systems also do not playback media. An example of a playback system that is not “media particular” would be an RF spectrum analyzer that records and plays back RF data. For example, an RF spectrum analyzer (e.g. <https://www.tek.com/application/rf-record-and-playback>) records and plays back RF data, not audio or video media. In my opinion, a POSITA could also understand “media particular playback system” as a subset of “playback systems.”

58. I disagree with Dr. Schmidt that a POSITA would understand the “media particular playback system” of claims 3, 15 or 26 to mean “media playback system.” I have reviewed the prosecution history, but find that it does not resolve the debate relating to the use of the term “particular.”

59. In my opinion the scope of this term is not reasonably ascertainable by those skilled in the art.

E. “data network” [’966, ’033, ’885 Patents]

Sonos Construction	Google Construction
“a medium that interconnects devices, enabling them to send digital data packets to and receive digital data packets from each other”	Plain and ordinary meaning; no construction necessary at this time

60. Sonos’s experts, Dr. Almeroth and Schmidt, opine that a POSITA would have understood the term “data network” to be limited to networks that transfer “digital data packets” and enable two-way communication. In the context of these patents, a person of ordinary skill in the art would have understood that the general understanding of the term “data network” does not restrict the type of data (digital or analog), the manner of transmission (packet or non-packet form), or the nature of the communication (bi-directional or unidirectional). Thus, a person of ordinary skill in the art would not understand “data network” to be restricted in the way Sonos is proposing, as I discuss in greater detail below.

F. Digital Data Packets

61. Sonos’ construction restricts the term “data network” to the transfer of “digital data packets.” In my opinion, a person of ordinary skill in the art would not understand a “data network” to be restricted to transferring “digital data packets.”

62. Numerous technical dictionaries confirm that “data”—including audio data—can be represented in both “analog” or “digital” form. Digital data is “[d]ata represented in discrete, discontinuous form, as contrasted with analog data represented in continuous form.” Ex. D (Modern Dictionary of Electronics, Seventh Edition (1999)). “Analog data” is “[d]ata represented in a continuous form, as contrasted with digital data represented in a discrete (discontinuous) form.” *Id.* (Modern Dictionary of Electronics, Seventh Edition (1999)). None of these dictionaries imply nor state that analog data may not be carried over a network. And relatedly, other publications also confirm that the plain and ordinary meaning of “data network” in the ’966, ’033 and ’995 patents encompasses both “digital” and “analog” data networks. *See, e.g.,* Ex. E (U.S. Patent Publication No. US2003/0087636) at ¶27 (“The embedded system may be capable of accessing various types of WANs, like a connection to a digital network or analog

data network.”); Ex. F (U.S. Patent No. 6,829,603) at 7:17-23 (“Such a session is maintained by a network interface 140 connecting to one or more of the following: the Internet 145, an intranet, a local area network, a public service telephone network, a wireless cellular network, a cable network, a satellite communications network or any other private or public digital or analog data network.”).

63. I also disagree with Dr. Almeroth and Schmidt that a data network is restricted to transferring only “packets.” “In the generic sense,” packets “refer[] to the manner in which data are organized into discrete units for transmission and switching through a data network.” Ex. G (Webster’s New World Telecom Dictionary (2008)). The plain and ordinary meaning of “data networks,” as understood by a person of skill in the art, would not be limited to only communicating data in digital data packets.

64. While some data networks transfer digital data packets, other do not. As one example, cellular networks were for many years analog rather than digital networks. Indeed, these networks were very popular and continuously running in the United States at the time that the Zone Scene patents were filed. <https://www.computerworld.com/article/2537122/most-analog-cellular-to-fade-away-on-monday.html>. These networks allowed cellular devices to send and receive data (as Sonos requires), typically in the form of voice calls. Likewise, a huge number of analog phone networks remain in the market, and for good reason: they are inexpensive, reliable, and have very high quality. <https://www.nojitter.com/sctc/analog-phones-no-plans-extinction>. Speaker system networks have also long used analog data. For example, RCA connections to speakers and other devices may deliver both video and audio data, and have been hugely popular. https://en.wikipedia.org/wiki/RCA_connector. Other types of non-

packetized data networks, such as circuit-switched data networks and public switched telephone network, were also widely used at the time the patents were filed.

G. Bi-Directional Data Networks

65. Sonos’ construction also restricts the term “data network” to a network that enables devices to send digital packets to and receive digital packets from each other. In other words, Sonos seeks to restrict the term “data network” to bi-directional data networks. In my opinion, a person of ordinary skill in the art would not understand a “data network” to be restricted to bi-directional data networks.

66. A person of ordinary skill in the art would understand the plain and ordinary meaning of “data network” encompasses both unidirectional and bidirectional data networks such that Dr. Almeroth and Schmidt’s contention that a network requires that a device must be able to send and receive data from another device is also incorrect. There are many types of networks that do not require a networked device to both send and receive data from another device. For example, networks may be configured in a ring, such that no device both sends and receives data directly to and from another device. *See* Ex. H (Webster’s New World Dictionary of Computer Terms) at 311. As the IEEE explained in its “IEEE Standard for Local and Metropolitan Area Networks: Overview and Architecture,” networks such as token-ring networks do not have the architecture described by Dr. Almeroth and Schmidt. Ex. I.

67. Various publications also confirm that “unidirectional” data networks were well-known in the art. For example, U.S. Patent No. 6,081,907 (“the ’907 patent”), which was filed on June 9, 1997, explains as part of its section on the state of the art that “[a]part from the classic bi-directional data networks,” there are an increasing number of “unidirectional” data networks, such as “broadcast or multicast networks.” ’907 patent at 1:36-47. The ’907 patent correctly

notes that these unidirectional networks support only “one-way communication,” such that “data flows from the server to the client,” but there is no “return communication path from the client to the server.” *Id.* at 1:36-59. The ’907 patent is directed at a system “in which computer data and other content are served from one or more servers over a unidirectional network to one or more clients.” *See, e.g.*, ’907 patent at 1:7-13, 1:64-2:8, 2:57-67, Claim 1. As another example, those of skill in the art recognized that networks may include “data diodes” (sometimes referred to as “unidirectional” gateways) to enforce data transfer in one direction between segments or devices of a network, for example to provide additional security for the network. Ex. G (Okhravi et al., Data Diodes in Support of Trustworthy Cyber Infrastructure) at § 2 (“Data diodes provide a physical mechanism for enforcing strict unidirectional communication between two networks.”), Fig. 2 (illustrating networks connected by data diode).

H. The Extrinsic Evidence Relied Upon By Dr. Almeroth And Schmidt Does Not Support Their Opinion

68. I have reviewed the extrinsic evidence cited by Drs. Almeroth and Schmidt in their declarations. In my opinion, when properly interpreted, the evidence does not support their opinions. Rather, it supports my opinion that a “data network” is not restricted in the manner Sonos proposes here.

69. For example, Dr. Almeroth cites to Data & Computer Communications (6th Ed. 2000) as support for this opinion. Almeroth Decl., ¶¶68-69 (emphasis added). Initially, this publication does not purport to define the term “data network.” Instead, it provides a general discussion of two “major categories into which communications networks are traditionally classified: wide area networks (WANs) and local area networks (LANs). Almeroth Decl., Appendix G at SONOS-SVG2-00018718. Even putting this fact aside, this publication actually

undercut Dr. Almeroth's opinions. For example, Dr. Schmidt opines that it teaches that "WANs" are "**generally** implemented using 'packet switching' and 'ATM' technologies." Schmidt Decl., ¶69 (emphasis added); *see also id.*, Appendix G at SONOS-SVG2-00018718 (explaining that in LANs data is "**usually**," but not always, transmitted in packets). Teaching that WANs "generally" use packets does not mean that "data networks" necessitate packets. To the contrary, it confirms that using the term "data networks" do not necessarily include packets. Moreover, Dr. Schmidt does not point to any teaching in the publication that indicates that data networks are not unidirectional.

70. In fact, Dr. Schmidt fails to mention that Data & Computer Communications (6th Ed. 2000) confirms that "Wide Area Networks" are not limited to "packet-switched networks." Schmidt Decl., Appendix G at SONOS-SVG2-00018719. Indeed, the publication identifies "circuit switching" and "telephone networks" as other examples of a "Wide Area Network." *Id.* As I explained above, circuit-switching and telephone networks may transfer data that is not in digital data packets.

71. Dr. Schmidt also points to several publications that discuss local area networks in the context of networks implementing the "OSI Reference Model" and "protocols such as IP." Schmidt Decl., ¶¶70-71; *see also* Schmidt Decl., Appendix C, D, H-I. A person of ordinary skill in the art would understand "OSI" and "IP" networks to refer to an exemplary type of "data network." Dr. Schmidt, however, does not dispute that the term "data network" is not restricted to a network implementing IP or the OSI model. In fact, as I showed above, Dr. Schmidt cites to publications which teach that the term "data network" encompasses other types of networks (*e.g.*, circuit-switched networks). A person of ordinary skill in the art would not understand the

description of a local area network in handbooks about packet-switched IP networks to provide a definition of other types of local area networks, or the more general term “data network.”

72. Dr. Schmidt also cites to the Webster’s New World Telecom Dictionary (2008) that describes “local area network” as “a packet network designed to interconnect host computers, peripherals, storage devices, and other computing resources within a local area, i.e., limited distance.”). Schmidt Decl., ¶75. In my experience, some definitions of technical terms may inadvertently limit the terms from their plain meanings. The existence of other publications cited by Dr. Schmidt, which I discussed above, confirm that a data network is not restricted to packets confirms that while data networks can be packet networks, not all data networks need to be packet networks.

73. Dr. Schmidt also cites to the Packet Broadband Network Handbook (2004). Schmidt Decl., ¶76. It is not surprising that this Handbook would describe packet networks since the Handbook is all about Packet Broadband Networks. A person of ordinary skill in the art would not understand the description of a local area network in a handbook about “packet broadband networks” to provide a definition of other types of local area networks, or the more general term “data network.”

I. “wherein the instruction comprises an instruction” [’033 Patent]

Sonos Construction	Google Construction
Plain and ordinary meaning; no construction necessary at this time	Indefinite

74. Claims 1 and 12 of the ’033 Patent recite “transmitting an instruction” and claims 2 and 13 recite “wherein the instruction comprises an instruction.” The specification of the ’033 Patent is largely silent with respect to the format or structure of instructions, and silent with respect to any instruction “compris[ing] an instruction.” *See e.g.* ’033 Patent at 3:1-7, 8:11-16.

Assuming “an instruction” means one instruction it is unclear to me, and would be unclear to a POSITA, where one “instruction” ends and another begins, or what constitutes one “instruction” as opposed to a plurality of “instructions.”

75. The specification of the '033 Patent is of no help to determine the meaning of the phrase “wherein the instruction comprises an instruction,” and it is my opinion that a person of ordinary skill in the art would be similarly unable to determine the meaning of the phrase. It would be unclear to a POSITA for example with respect to a block of code, which parts of the block of code formed “the instruction” itself, and which parts of the block would be “an instruction.” It would be similarly unclear to a POSITA what the exact bounds of the block of code or “the instruction” would be, i.e. where one instruction ends and another begins. Several lines of codes could be understood as one instruction or five instructions. For example, is one line of code an “instruction”? One function or method? Is one “for” loop within such a method one “instruction” or many “instructions”? And how does POSITA deal with method calls within other method calls? Looking at lower and high levels in the stack adds further difficulty to this question—would a line of assembly code constitute one instruction?

76. Accordingly, it is my opinion that this term is indefinite because a POSITA could not determine with certainty the metes and bounds of the term.

DATED: June 1, 2021

Respectfully submitted,

By /s/ Christos Kyriakakis

Christos Kyriakakis

EXHIBIT A

Prof. Chris Kyriakakis

Contact

USC Viterbi School of Engineering (1996 – Present)
 Ming Hsieh Department of Electrical Engineering
 email: ckyriak@me.com
 Wikipedia: http://en.wikipedia.org/wiki/Chris_Kyriakakis

Education

PhD 1993	University of Southern California (USC) Electrical Engineering
MS 1987	University of Southern California (USC) Electrical Engineering
BS 1985	California Institute of Technology (Caltech) Engineering and Applied Science

Positions

Associate Professor 2002-Present	USC Electrical Engineering
Assistant Professor 1996-2002	USC Electrical Engineering
Research Fellow 1993-1996	USC Electrical Engineering
CTO 2002-2018	Audyssey Laboratories
Chief Audio Scientist 2019-Present	SYNG

Research Areas

Immersive audio signal processing, machine learning for audio and acoustics, room equalization, room acoustics, headphone acoustics, multichannel audio streaming, microphone array processing, head-related transfer function modeling, adaptive algorithms for virtual microphones, virtual loudspeakers, multichannel audio rendering, archaeoacoustics.

Honors and Awards

2018	Elected as Senior Member of the Institute of Electrical and Electronics Engineers (IEEE)
2006	World Technology Network Award

James P. Clark, Chairman of the World Technology Network: “The World Technology Awards program is a very inspiring way to identify and honor the most innovative people and organizations in the technology world...We look forward to assisting Prof. Kyriakakis continue to help create our collective future and change our world.”

- 2004 National Academy of Engineering: Frontiers of Engineering Session Chairman
- 2003 Best Paper Award: IEEE 37th Asilomar Conference on Signals, Systems, & Computers

Professional Service (Past and Current)

Society Memberships

Member of the Audio Engineering Society (AES)
Senior Member of the IEEE
Member of the Acoustical Society of America (ASA)
Member IEEE Technical Committee on Multimedia
Member of AES Technical Committee on Multichannel and Binaural Audio
Member of AES Technical Committee on Audio for Games

Reviewer

National Science Foundation
IEEE Transactions on Multimedia
IEEE Transactions on Signal Processing
IEEE Transactions on Audio, Speech, and Language Processing
European Association for Signal Processing (EURASIP)
IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)
IEEE Workshop on Applications of Signal Processing to Audio (WASPAA)
IEEE Workshop on Multimedia Signal Processing
ACM Special Interest Group on Graphics and Interactive Techniques (SIGGRAPH)
European Signal Processing Conference (EUSIPCO)
Audio Engineering Society
Acoustical Society of America

Books

1. D. Yang, C.-C. Kuo, and C. Kyriakakis, "*High Fidelity Multichannel Audio Coding*," Hindawi Press, New York, (2004).
2. S. Bharitkar and C. Kyriakakis, "*Immersive Audio Signal Processing*," Springer-Verlag, New York, (2006).

Patents

1. Audio Output Device Specific Audio Processing. 9,860,641 (2018)
2. Audio Content Enhancement Using Bandwidth Extension Techniques. 8,705,764 (2014)
3. Room Acoustic Response Modeling and Equalization with Linear Predictive Coding and Parametric Filters. 8,363,853 (2013)
4. Cross-Over Frequency Selection and Optimization of Response Around Crossover. 8,363,852 (2013)
5. Stereo Expansion with Binaural Modeling. 8,229,143 (2012)
6. Phase Equalization for Multi-Channel Loudspeaker-Room Responses. 8,218,789 (2010)
7. Combined Multirate-Based and FIR-Based Filtering Technique for Room Acoustic Equalization. 8,077,880 (2011)
8. System and Method for Automatic Multiple Listener Room Correction with Low Order Filters. 8,005,228 (2011)
9. Personal Audio Docking Station. D629,391 (2010)
10. Cross-Over Frequency Selection and Optimization of Response Around Cross-Over. 7,826,626 (2010)
11. System and Method for Automatic Room Acoustic Correction in Multi-Channel Audio Environments. 7,769,183 (2010)
12. Phase Equalization for Multi-Channel Loudspeaker-Room Responses. 7,720,237 (2010)
13. System and Method for Automatic Multiple Listener Room Acoustic Correction with Low Filter Orders. 7,567,675 (2009)

Selected Press Coverage

December 2, 2017	BBC World Service: Symphony of the Stones http://www.bbc.co.uk/programmes/w3csvq4j
February 28, 2017	National Herald: Professors Recreate the Sounds of an Ancient Greek Church in LA https://www.thenationalherald.com/152455/professors-recreate-sounds-ancient-greek-church-la/
February 19, 2016	The Atlantic: Hearing the Lost Sounds of Antiquity http://www.theatlantic.com/technology/archive/2016/02/byzantine-angel-wings/470076/
September 14, 2015	UCLA Newsroom: Measuring the sound of angels singing http://newsroom.ucla.edu/stories/measuring-the-sound-of-angels-singing
February 5, 2015	Gizmodo: No, women's voices are not easier to understand than men's voices http://gizmodo.com/no-siri-is-not-female-because-womens-voices-are-easier-1683901643
September 30, 2014	IEEE Spectrum: Why mobile voice quality still stinks and how to fix it http://spectrum.ieee.org/telecom/wireless/why-mobile-voice-quality-still-stinksand-how-to-fix-it
June 21, 2012	NPR All Things Considered: He'll Retune Your Living Room http://www.npr.org/blogs/therecord/2012/06/21/155501975/hell-retune-your-living-room
Sept. 5, 2011	New York Times: Sound, the Way the Brain Prefers to hear it http://www.nytimes.com/2011/09/06/science/06sound.html?_r=1&pagewanted=all
Nov. 29, 2006	Widescreen Review: Chris Kyriakakis wins 2006 World Technology Award http://www.widescreenreview.com/news_detail.php?id=12381
Oct. 27, 2006	USC News: Immersive Sound-The Next Step in Audio http://www.usc.edu/uscnnews/stories/12915.html

Journal Publications

1. S. Gerstel, C. Kyriakakis, S. Antonopoulos, K. Raptis, and J. Donahue, "Soundscapes of Byzantium: The Acheiropoietos Basilica and the Cathedral of Hagia Sophia in Thessaloniki," *Hesperia*, Issue 87.1, January-March 2018
2. S. Gerstel, C. Kyriakakis, S. Antonopoulos, K. Raptis, and J. Donahue, "Soundscapes of Byzantium", *Speculum* 92(5), October 2017
3. S. Gerstel and C. Kyriakakis, "Uncovering the Acoustic Mysteries of Byzantine Thessaloniki", *Faith and Form*, 49(3), Oct. 2016
4. C. S. Lin and C. Kyriakakis, "High Definition Audio Spectrum Separation via Frequency Response Masking Filter Banks", *Journal of the Audio Engineering Society*, 57(10), 807-821, (2009)
5. C. S. Lin and C. Kyriakakis, "Removing Additive Noise via Neuro-Fuzzy Based Reinforcement Learning," *Journal of the Acoustical Society of America*, 124(2), 1026-1037, (2008)
6. D. Cantzos and C. Kyriakakis, "Quality Enhancement of Compressed Audio Based on Statistical Conversion," *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2008, p. 1-15, (May 2008).
7. S. Bharitkar and C. Kyriakakis, "Visualization of Multiple Listener Room Acoustic Equalization with Sammon Map," *IEEE Transactions on Speech and Audio Processing*, 15(2), 542-551, (2007).
8. P. G. Georgiou and C. Kyriakakis, "Robust Maximum Likelihood Source Localization: The Case for Sub-Gaussian versus Gaussian," *IEEE Transactions on Speech and Audio Processing*, 14(4), 1470-1480, (2006).
9. P. G. Georgiou and C. Kyriakakis, "Maximum Likelihood Parameter Estimation under Impulsive Conditions, a Sub-Gaussian Signal Approach," *EURASIP Signal Processing Journal*, 10, 3061-3075 (2006).
10. A. Mouchtaris, S. S. Narayanan, and C. Kyriakakis, "Multichannel Audio Synthesis by Subband-Based Spectral Conversion and Parameter Adaptation," *IEEE Trans. Speech and Audio Processing*, 13(2), 263-274, (2005).

11. S. Bharitkar and C. Kyriakakis, "Robustness of Spatial Average Equalization: A Statistical Reverberation Model Approach," *Journal of the Acoustical Society of America*, 116(6), 3491-3497, (2004).
11. R. Zimmermann, C. Kyriakakis, C. Shahabi, C. Papadopoulos, A. A. Sawchuk, U. Neumann, "The Remote Media Immersion System: Internet Meets the Future Home Theater," *IEEE Multimedia Magazine*, March 2004.
12. S. Bharitkar and C. Kyriakakis, "Selective Signal Cancellation For Multiple-Listener Audio Applications Using Eigenfilters," *IEEE Transactions on Multimedia*, vol. 5(3), 329-338, (2003).
13. D. Yang, H. Ai, C. Kyriakakis, and C. -C. Jay Kuo, "High-Fidelity Multichannel Audio Coding With Karhunen-Loeve Transform," *IEEE Transactions on Speech and Audio Processing*, 11(4), 365-380, (2003).
14. A. Mouchtaris, S. S. Narayanan, and C. Kyriakakis, "Virtual Microphones for Multichannel Audio Resynthesis," *EURASIP Journal on Applied Signal Processing*, Special Issue on Digital Audio for Multimedia Communications, 10, 968-979, (2003).
15. D. Yang, H. Ai, C. Kyriakakis and C.-C. Jay Kuo, "Progressive Syntax-Rich Audio Coding for Multichannel Audio Sources," *EURASIP Journal on Applied Signal Processing*, Special Issue on Digital Audio for Multimedia Communications, 10, 980-992, (2003).
16. A. Mouchtaris, P. Reveliotis, and C. Kyriakakis, "Inverse Filter Design for Immersive Audio Rendering over Loudspeakers," *IEEE Transactions on Multimedia*, 2(2), 77-87, (2000).
17. P. Georgiou, C. Kyriakakis, and P. Tsakalides, "Alpha Stable Modeling of Noise and Robust Time-Delay Estimation in the Presence of Impulsive Noise," *IEEE Transactions on Multimedia*, 1(3), 291-301, (1999).
18. C. Kyriakakis, P. Tsakalides, and T. Holman, "Surrounded by Sound: Immersive Audio Acquisition and Rendering Methods," *IEEE Signal Processing Magazine*, 16(1), 55-66, (1999).
19. C. Kyriakakis, T. Holman, J.-S. Lim, H. Hong, and H. Neven, "Signal Processing, Acoustics, and Psychoacoustics for High Quality Desktop Audio," *Journal of Visual Communication and Image Representation*, 9(1), 51-61, (1998).

20. C. Kyriakakis, "Fundamental and Technological Limitations of Immersive Audio Systems," Proceedings of the IEEE, **86**(5), 941-951, (1998).

Technical Conference Publications

1. D. Cantzos, A. Mouchtaris, and C. Kyriakakis, "Perceptually-Driven Scalable MDCT Enhancement of Compressed Audio Based on Statistical Conversion", Proceedings of IEEE International Symposium on Multimedia, Dec. 2011.
2. S. Bharitkar and C. Kyriakakis, "Comparison between Spherical Headmodels and HRTFs in Upmixing for Headphone Based Virtual Surround and Stereo Expansion", Proceedings of the AES 129th Convention in San Francisco, CA, November 4-7, 2010.
3. C. S. Lin, "Design of Phase Equalizer for Rational Filter Banks Using Minimax Approach", IEEE International Symposium on Industrial Electronics, June 2009.
4. D. Cantzos and C. Kyriakakis, "Bandwidth Extension of Low Bitrate Compressed Audio Based on Statistical Conversion", Proceedings of IEEE ICME 2009.
5. S. Bharitkar, C. Kyriakakis, T. Holman, "Variable-Octave Complex Smoothing for Loudspeaker-Room Response Equalization", IEEE International Conference on Consumer Electronics (ICCE 2008), Las Vegas, NV, 9-13 Jan., 2008.
6. D. Cantzos, A. Mouchtaris, and C. Kyriakakis, "Synthesis of Enhanced Audio from Low Bitrate Compressed Audio Based on Unit Selection and Statistical Conversion Methods," 42nd Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, CA, Oct. 26, 2008.
7. S. Bharitkar, Y. Zhang, and C. Kyriakakis, "Room Acoustic Response Modeling and Equalization with Linear Predictive Coding and Parametric Filters for Speech and Audio Enhancement," in *Proc. 40th IEEE Asilomar Conf. on Signals, Systems, & Computers*, Pacific Grove (CA), Oct. 2006.
8. S. Bharitkar and C. Kyriakakis, "Cascaded FIR Filters for Multiple Listener Low Frequency Room Acoustic Equalization," accepted in *2006 IEEE Intl. Conf. Acoust.*,

- Speech, & Signal Proc. (ICASSP-2006)*, Toulouse, France, May 2006.
9. C. -S. Lin and C. Kyriakakis, "A Morphing Approach for Synthesizing Multichannel Recordings," 39th IEEE Asilomar Conference on Signals, Systems, & Computers, Pacific Grove, CA, November, 2005.
 10. J. Peterson and C. Kyriakakis, "Analysis of Fast Localization Algorithms for Acoustical Environments," 39th IEEE Asilomar Conference on Signals, Systems, & Computers, Pacific Grove, CA, November, 2005.
 11. S. Bharitkar and C. Kyriakakis, "Comparison Between Time Delay Based and Nonuniform Phase Based Equalization for Multichannel Loudspeaker-Room Responses," 39th IEEE Asilomar Conference on Signals, Systems, & Computers, Pacific Grove, CA, November, 2005.
 12. S. Bharitkar and C. Kyriakakis, "Optimization of the Bass Management Filter Parameters for Multichannel Audio Applications," 39th IEEE Asilomar Conference on Signals, Systems, & Computers, Pacific Grove, CA, November, 2005.
 13. S. Bharitkar and C. Kyriakakis, "Objective Function for Automatic Multiposition Equalization and Bass Management Filter Selection," Audio Engineering Society, New York, October, 2005.
 14. S. Sundaram and C. Kyriakakis, "Phantom Audio Sources With Vertically Separated Speakers," Audio Engineering Society, New York, October, 2005.
 15. S. Bharitkar and C. Kyriakakis, "Optimization of Crossover Frequency and Crossover Region Response for Multichannel Acoustic Applications," 13th European Signal Processing Conference (EUSIPCO 2005), Antalya, Turkey, September, 2005.
 16. J. Peterson and C. Kyriakakis, "Hybrid Algorithm for Robust, Real-Time Source Localization in Reverberant Environments," International Conference on Acoustics, Speech, and Signal Processing (ICASSP '05), Philadelphia, March 2005.
 17. J. Peterson and C. Kyriakakis, "Choosing Candidate Locations for Source Localization," Joint Workshop on Hands-Free Speech Communication and Microphone Arrays (HSCMA '05), Rutgers, March 2005.

18. C. Papadopoulos, C. Kyriakakis, A. A. Sawchuk, X. He, "Cyberseer: 3D Audio-Visual Immersion for Network Security and Management," Workshop on Visualization and Data Mining for Computer Security (VizSEC/DMSEC) 2004, Fairfax, VA, October, 2004.
19. H. Huang and C. Kyriakakis, "Binaural Noise Reduction Combining Binaural Analysis and Spectral Subtraction," Proc. 38th IEEE Asilomar Conference on Signals, Systems, & Computers, Pacific Grove, CA. Nov. 2004.
20. H. Huang and C. Kyriakakis, "Computationally Efficient Blind Dereverberation of Audio Signals," Proceedings of the 117th Convention of the Audio Engineering Society, New York, October, 2004.
21. M. Boueri and C. Kyriakakis, "Audio Signal Decorrelation based on a Critical Band Approach," Proceedings of the 117th Convention of the Audio Engineering Society, New York, October, 2004.
22. S. Bharitkar and C. Kyriakakis, "Nonlinear Projection Algorithm for Evaluating Multiple Listener Equalization Algorithm Performance," Proceedings of the 117th Convention of the Audio Engineering Society, New York, October, 2004.
23. R. Sadek and C. Kyriakakis, "A Novel Multichannel Panning Algorithm for Standard and Arbitrary Loudspeaker Configurations," Proceedings of the 117th Convention of the Audio Engineering Society, New York, October, 2004.
24. A.A. Sawchuk, E. Chew, R. Zimmermann, C. Papadopoulos and C. Kyriakakis, "Distributed Immersive Performance," Corporation for Education Network Initiatives in California (CENIC) Eighth Annual Conference, March 15-17, 2004, Marina del Rey, CA.
25. C. S. Lin and C. Kyriakakis, "Multi-Frequency Noise Removal Based on Reinforcement Learning," Preprint No. 5966, presented at the 115th AES Convention, New York, NY, October, 2003.
26. S. Bharitkar, P. Hilmes, C. Kyriakakis, "Sensitivity of Multichannel Room Equalization to Listener Position," in Proceedings of the AES 115th Convention, NY, Oct 2003.
27. H. Huang and C. Kyriakakis, "Subband Adaptive Filtering For Acoustic Noise Reduction," in Proceedings of the AES 115th Convention, NY, Oct 2003.

28. S. Bharitkar and C. Kyriakakis, "The Influence of Reverberation on Multichannel Equalization: An Experimental Comparison Between Methods," Proc. 37th IEEE Asilomar Conference on Signals, Systems, & Computers, Pacific Grove, CA. Nov. 2003. **(Best Paper Award)**
29. J. M. Peterson, P. Georgiou, and C. Kyriakakis, "The Use of Fractional Lower-Order Statistics in acoustical environments," in Proceedings of Asilomar Conference on Signals, Systems, and Computers, 2003.
30. C. S. Lin and C. Kyriakakis, "Frequency Response Masking Approach for Designing Filter Banks with Rational Sampling Factors," IEEE WASPAA 2003, Mohonk, NY, October, 2003.
31. S. Bharitkar and C. Kyriakakis, "A Comparison Between Multi-channel Audio Equalization Filters Using Warping," 2003 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA 2003), New Paltz, NY. Oct. 2003.
32. S. Bharitkar and C. Kyriakakis, "Sensitivity of multichannel room equalization to listener position," Proc. IEEE Conf. on Multimedia & Expo (ICME), Baltimore. July 2003.
33. S. Bharitkar, P. Hilmes, C. Kyriakakis, "Analysis of Multiple Listener Equalization Performance Due to Displacement Effects," 145th Meeting of the Acoustical Soc. America. Nashville (TN). May 2003.
34. A.A. Sawchuk, C. Kyriakakis, R. Zimmermann, C. Papadopoulos, C. Shahabi, U. Neumann and T. Holman, "Remote Media Immersion," Corporation for Education Network Initiatives in California (CENIC) Seventh Annual Conference, May 7-9, 2003, Santa Barbara, CA.
35. S. Rishi, C. Papadopoulos, and C. Kyriakakis, "Loss Concealment for Multichannel Streaming Audio," Network and Operating System Support for Digital Audio and Video (NOSSDAV '03), Monterey, CA, June, 2003.
36. P. G. Georgiou and C. Kyriakakis, "A new model of source dependent noise for robust array signal processing," in Second IEEE Sensor Array and Multichannel Signal Processing Workshop, Aug. 2002.
37. P. G. Georgiou and C. Kyriakakis, "Signal Parameter Estimation and Localization via Maximum Likelihood

- Using and Sensor Array in the Presence of Levy Noise,” ICSP 2002, Beijing.
38. S. Bharitkar and C. Kyriakakis, “Perceptual Multiple Location Equalization with Clustering,” in Proc. 36th IEEE Asilomar Conference on Signals, Systems, & Computers, Pacific Grove, CA, Nov. 2002.
 39. S. Bharitkar, P. Hilmes, and C. Kyriakakis, “Robustness of Spatial Averaging Equalization Methods: A Statistical Approach,” in Proc. 36th IEEE Asilomar Conference on Signals, Systems, & Computers, Pacific Grove, CA, Nov. 2002.
 40. G. Georgiou and C. Kyriakakis, “A novel model for reverberant signals; robust maximum likelihood localization of real signals based on a Sub-Gaussian model,” in Proc. 36th IEEE Asilomar Conference on Signals, Systems, & Computers, Pacific Grove, CA, Nov. 2002.
 41. A. Mouchtaris, S. S. Narayanan, and C. Kyriakakis, “Maximum Likelihood Constrained Adaptation for Multichannel Audio Synthesis,” in Proc. 36th IEEE Asilomar Conference on Signals, Systems, & Computers, Pacific Grove, CA, Nov. 2002.
 42. S. Bharitkar and C. Kyriakakis, “Visualization of Multiple Listener Room Response Equalization using Sammon Map,” in Proc. IEEE Intl. Conf. on Multimedia & Expo (ICME 2002), Lausanne, Switzerland, Aug. 2002.
 43. Mouchtaris, S. S. Narayanan, and C. Kyriakakis, “Multiresolution Spectral Conversion for Multichannel Audio Resynthesis,” in Proc. IEEE International Conference on Multimedia and Expo (ICME 2002), vol. 2, p. 273-276, Lausanne, Switzerland, Aug. 2002.
 44. A. Mouchtaris, S. S. Narayanan, and C. Kyriakakis, “Efficient Multichannel Audio Resynthesis by Subband-Based Spectral Conversion,” in Proc. European Processing Conference (EUSIPCO), vol. 1, p. 413-416, Toulouse, France, Sept. 2002.
 45. P. G. Georgiou and C. Kyriakakis, “A robust array signal processing maximum likelihood estimator based on sub-gaussian signals,” in Proc. European Processing Conference (EUSIPCO), Toulouse, France, Sept. 2002.
 46. Yang, H. Ai, C. Kyriakakis, and C. -C. Jay Kuo, “Design of Progressive Syntax-Rich Multichannel Audio Codec,”

- Proc. SPIE Vol. 4674, p. 121-123, Media Processors 2002, San Jose, CA, January, 2002.
47. Yang, H. Ai, C. Kyriakakis and C.-C. Jay Kuo, "Error-Resilient Design of High Fidelity Scalable Audio Coding," Proc. SPIE Vol. 4740, Digital Wireless Communications IV, Orlando, Florida USA, April 1-5, 2002.
 48. J. M. Peterson and C. Kyriakakis, "Using subband filters to reduce the complexity of real-time signal processing," Preprint No. 5663, presented at the 113th AES Convention, Los Angeles, CA, October, 2002.
 49. D. Yang, H. Ai, C. Kyriakakis, and C. -C. Jay Kuo, "Design of Error-Resilient Layered Audio Codec," Preprint No. 5681, presented at the 113th AES Convention, Los Angeles, CA, October, 2002.
 50. S. Bharitkar, P. Hilmes, and C. Kyriakakis, "A Statistical Analysis of the Robustness of Room Equalization Using Spatial Averaging," Preprint No. 5669, presented at the 113th AES Convention, Los Angeles, CA, October, 2002.
 51. P. G. Georgiou and C. Kyriakakis, "A Novel Method for Source Localization using Microphone Arrays in Reverberant Environments," Preprint No. 5643, presented at the 113th AES Convention, Los Angeles, CA, October, 2002.
 52. A. Mouchtaris, S. S. Narayanan, and C. Kyriakakis, "Gaussian-Mixture-Model-Based Methods for Multichannel Audio Resynthesis," Preprint No. 5647, presented at the 113th AES Convention, Los Angeles, CA, October, 2002.
 53. C.-S. Lin and C. Kyriakakis, "A Fuzzy Cerebellar Model Approach for Synthesizing Multichannel Recordings," Preprint No. 5675, presented at the 113th AES Convention, Los Angeles, CA, October, 2002.
 54. S. Bharitkar and C. Kyriakakis, "A Cluster Centroid Method for Room Response Equalization at Multiple Locations," IEEE WASPAA 2001, Mohonk, NY (WASPAA: Workshop on Applications of Signal Processing to Audio and Acoustics).
 55. J. -S. Lim and C. Kyriakakis, "Multirate Adaptive Filtering for Immersive Audio," ICASSP '01, Salt Lake City, Utah, May, 2001 (ICASSP: International Conference on Acoustics, Speech, and Signal Processing).

56. Yang, H. Ai, C. Kyriakakis, and C. –C. Kuo, “Adaptive Karhunen-Loeve Transform for Enhanced Multichannel Audio Coding, SPIE, San Diego, 2001.
57. Yang, H. Ai, C. Kyriakakis, and C. –C. Kuo, “Embedded High-Quality Multichannel Audio Coding,” Conference on Media Processors, Symposium on Electronic Imaging 2001, San Jose, CA, January 21-26, 2001.
58. S. Bharitkar and C. Kyriakakis, “A Classification Scheme for Room Acoustical Responses,” ICASSP 2001, Salt Lake City, Utah.
59. J. –S. Lim and C. Kyriakakis, “Adaptive Synthesis of Immersive Audio Rendering Filters,” 111th AES Convention, New York, NY, 2001.
60. S. Bharitkar and C. Kyriakakis, “New Factors in Room Equalization Using a Fuzzy Logic Approach, , 111th AES Convention, New York, NY, 2001.
61. A. Mouchtaris and C. Kyriakakis, “Time Frequency Methods for Virtual Microphone Signal Synthesis,” 111th AES Convention, New York, NY, 2001.
62. S. Bharitkar and C. Kyriakakis, “Eigenfilters for Signal Cancellation,” ISPACS 2000, Hawaii.
63. C. Kyriakakis and A. Mouchtaris, “Virtual Microphones for Multichannel Audio Applications,” ICME 2000, New York, July 2000.
64. S. Bharitkar and C. Kyriakakis, “Selective Signal Cancellation for Multiple Listener Audio Applications: An Information Theory Approach,” ICME 2000, New York, July 2000.
65. P. Georgiou and C. Kyriakakis, “A Multiple Input Single Output Model for Immersive Audio Rendering in Real Time,” ICME 2000, New York, July 2000.
66. S. Rizzo, J. G. Buckwalter, L. Humphrey, C. van der Zaag, T. Bowerly, C. Chua, U. Neumann, and C. Kyriakakis, “A Virtual Reality Classroom Scenario for ADHD Assessment,” The 28th Annual Conference of the International Neuropsychological Society, Denver, Colorado, February, 2000.
67. P. G. Georgiou and C. Kyriakakis, “Modeling of Head-Related Transfer Functions for Immersive Audio Using a State-Space Approach,” 33rd Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, California, October, 1999.

68. A. Mouchtaris, Z. Zhu, and C. Kyriakakis, "High-Quality Internet Audio over ATM Networks," 33rd Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, California, October, 1999.
69. C. Kyriakakis, S. Zeadally, and T. Holman, "Multichannel Audio over Internet2 Using ATM," Audio Engineering Society, preprint no. 5000, (1999).
70. C. Kyriakakis, "Immersive Audio Rendering on the TI C62 DSP Platform," TI DSPFest, Houston, TX, August, 1999
71. P. G. Georgiou, P. Tsakalides, and C. Kyriakakis, "Alpha-Stable Robust Modeling of Background Noise for Enhanced Sound Source Localization," ICASSP '99, Phoenix, Arizona, March, 1999.
72. A. Mouchtaris, P. Reveliotis, and C. Kyriakakis, "Non-minimum Phase Inverse Filtering Methods for Immersive Audio Rendering," ICASSP '99, Phoenix, Arizona, March, 1999.
73. C. Kyriakakis and T. Holman, "Video-Based Head Tracking for Improvements in Multichannel Loudspeaker Audio," Audio Engineering Society, preprint no. 4845, (1998).
74. A. Mouchtaris, J.-S. Lim, T. Holman, and C. Kyriakakis, "Head-Related Transfer Function Synthesis for Immersive Audio," Proceedings of the IEEE Multimedia Signal Processing Workshop (MMSP '98), Redondo Beach, California, December, 1998.
75. C. Kyriakakis and T. Holman, "Immersive Audio for Desktop Systems," Acoustical Society of America International Conference on Acoustics, Seattle Washington, June, 1998
76. Kyriakakis and T. Holman, "Immersive Audio for the Desktop: Signal Processing Requirements," ICASSP '98, Seattle, Washington, May, 1998.
77. Mouchtaris, J.-S. Lim, T. Holman, and C. Kyriakakis, "Signal Processing Considerations for Immersive Audio Rendering," Proceedings of the 10th Tyrrhenian Conference on Multimedia Communications, Ischia, Italy, (1998).
78. P. Georgiou, C. Kyriakakis, and P. Tsakalides "Microphone Arrays for Speaker Direction Identification," Proceedings of the IEEE on Applications of Signal Processing to Audio and Acoustics, WASPAA, (1997).

79. T. Holman and C. Kyriakakis, "Acoustics and Psychoacoustics of Desktop Systems," 134th meeting of the Acoustical Society of America, San Diego, CA, (1997).

EXHIBIT B

Kyriakakis Materials Considered**Patents & File Histories**

U.S. Patent No. 9,219,460
U.S. Patent No. 9,344,206
U.S. Patent No. 9,967,615
U.S. Patent No. 10,469,966
U.S. Patent No. 10,779,033
U.S. Patent No. 9,219,460 File History
U.S. Patent No. 9,344,206 File History
U.S. Patent No. 9,967,615 File History
U.S. Patent No. 10,469,966 File History
U.S. Patent No. 10,779,033 File History

P.L.R. Exchanges

Sonos' Opening Claim Construction Brief and attachment/exhibits thereto (DN 60, filed 4/27/2021)
Google's Proposed Claim Terms For Construction (served: 2021.03.18)
Sonos' Proposed Claim Terms For Construction (served: 2021.03.18)
Google's Identification of Preliminary Claim Constructions (served: 2021.04.02)
Sonos' CORRECTED List of Proposed Constructions (served: 2021.04.02)
Google's Revised Identification of Preliminary Claim Constructions (served: 2021.04.16)

Productions/Extrinsic Evidence

Google's Disclosures of Extrinsic Evidence and attachments/exhibits thereto (served: 2021.04.09)
Sonos' Identification of Extrinsic Evidence and attachments/exhibits thereto (served 2021.04.09)

EXHIBIT C



THE UNIVERSITY OF BRITISH COLUMBIA

Deploying the world's largest campus IEEE 802.11b network

IEEE 802.11 November 2003 Plenary

Tutorial #5 – Case Study

Tuesday, November 11, 2003

Albuquerque Convention Center – Ballroom A

Jonn Martell, Wireless Project and Service Manager

University of British Columbia

jonn.martell@ubc.ca



THE UNIVERSITY OF BRITISH COLUMBIA

About the presentation and presenter

- Implementation of the world's largest IEEE 802.11b network.
- Showcase of IEEE standards in action.
- Provide feedback to IEEE members on where implementers need help.
- Jonn Martell
 - 15 years experience in network implementation.
 - 5+ years in wireless networking
 - IEEE 802.11 member



THE UNIVERSITY OF BRITISH COLUMBIA

Agenda

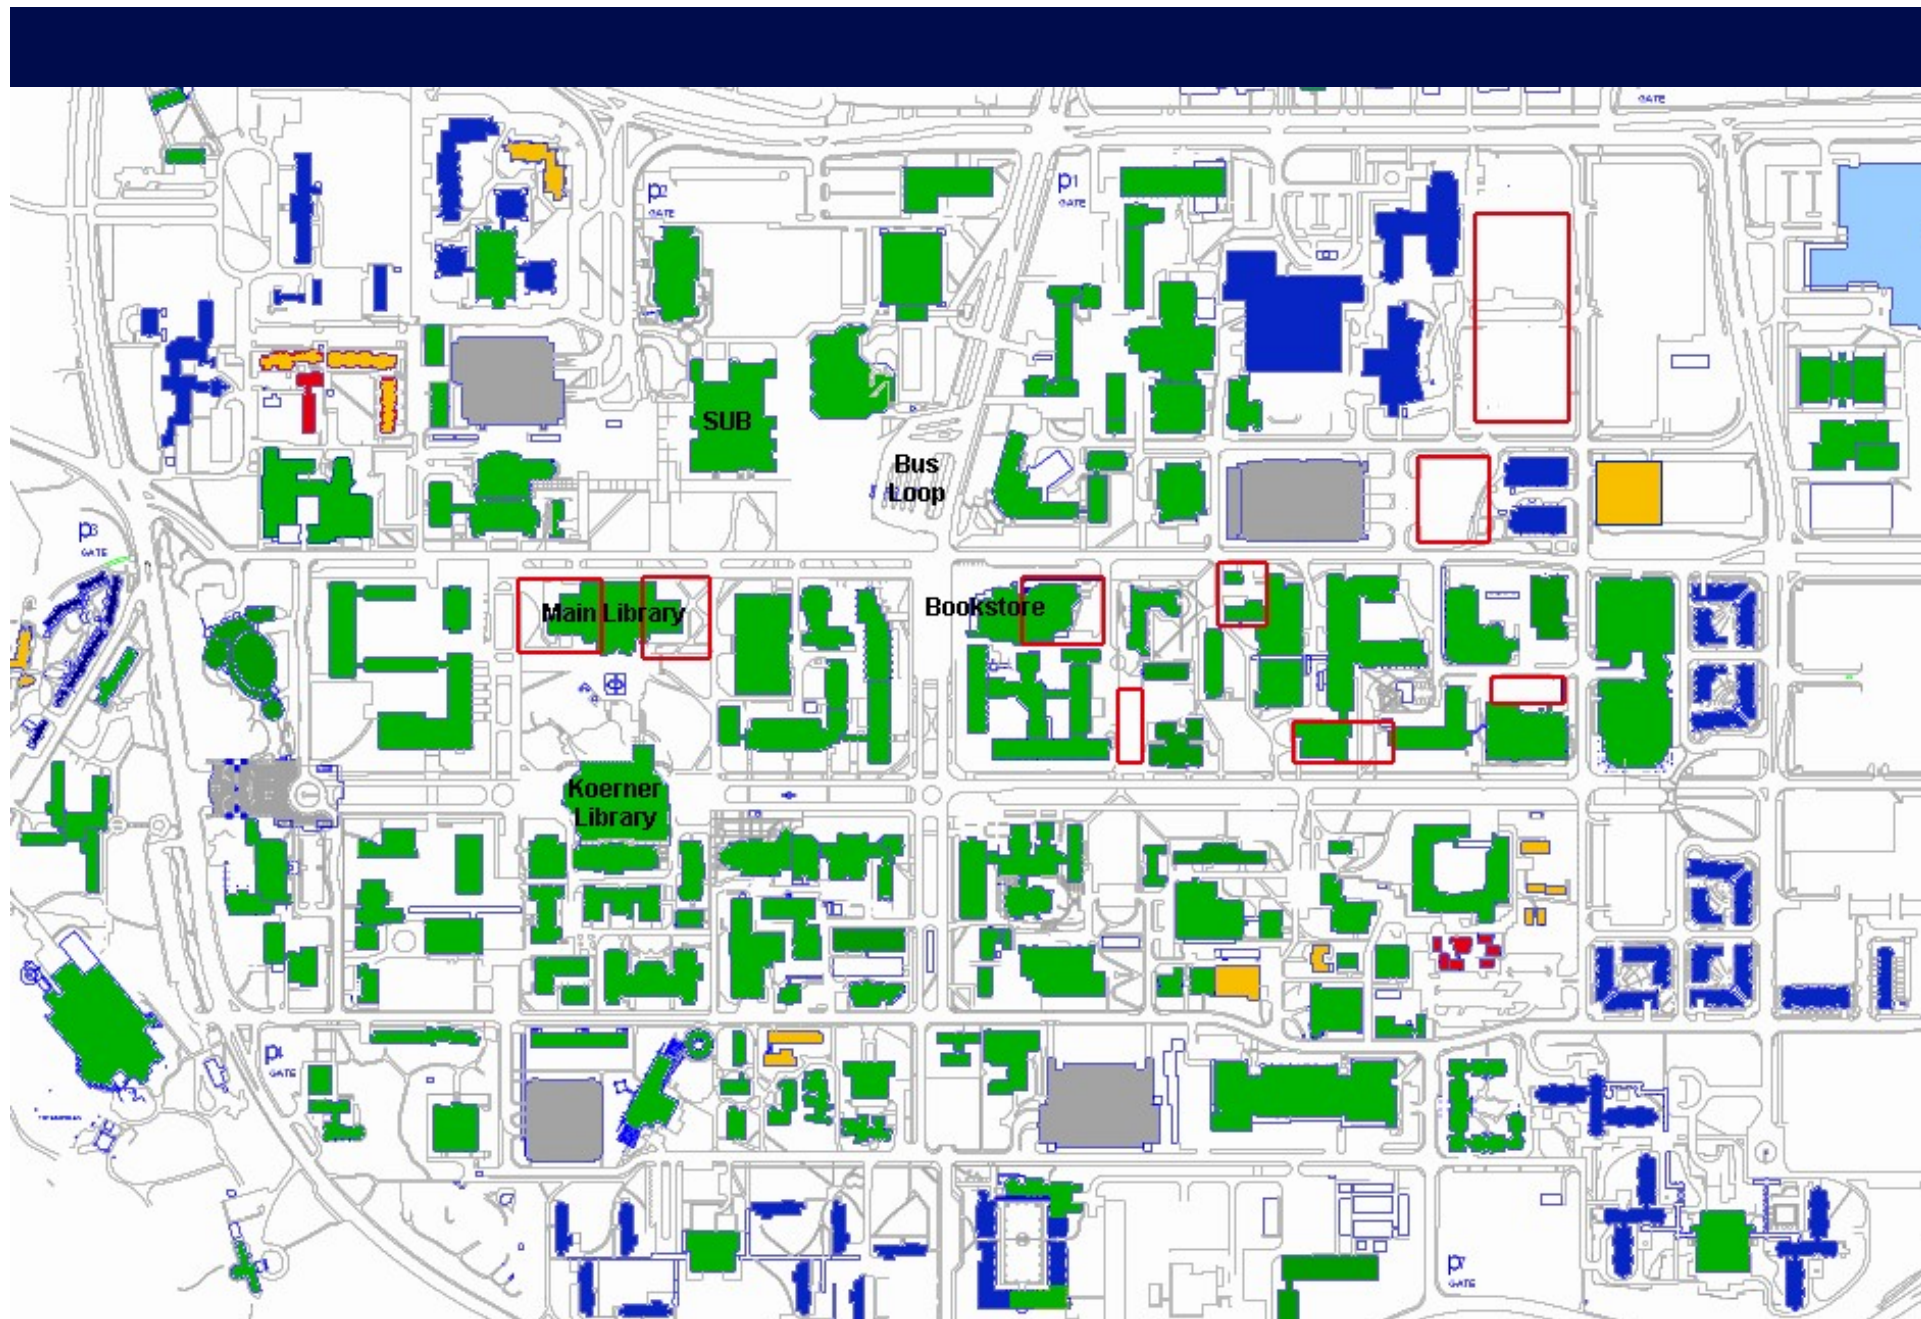
- Snapshot of wireless.ubc.ca
- Education environments
- Seamless campus-wide wireless network
- End user experience
- Wireless statistics and usage
- Challenges
- Key success factors
- Network design
- Security
- Management
- Futures and managed spectrum



THE UNIVERSITY OF BRITISH COLUMBIA

Snapshot of wireless.ubc.ca

- Campus wide
- Close to 5000 users
- 150 buildings
- 1300 Access Points (APs)
- 600 acres coverage
- 5 million square feet of coverage
- Roaming enabled
- Complete indoor and outdoor coverage
- \$5.9M CDN of a \$30.6M connectivity project
- Main campus completed (on time and on budget). Adding student residences.





THE UNIVERSITY OF BRITISH COLUMBIA

Education environments - EDUs

- Usability is more important than security.
- Mix of hotspot/public access and “enterprise” networking.
- Four different types of users.
- Relatively insecure indoor environments.
- Very decentralised; leadership is only valued by addressing user needs.
- High intellectual capability and autonomy.
- Will always need wired connectivity for high bandwidth applications (like video and medical applications).



THE UNIVERSITY OF BRITISH COLUMBIA

Motivation for a seamless campus wide wireless network

- There was a need, even two years ago!
- Choice: Deploy as a campus wide network or deal with hundreds of poorly configured APs & incompatible domains.
- Primary target, students and faculty are mobile and need campus wide access.
- Need to be ready for future applications like mobile phones.
- Need to avoid re-authentication issues between zones.
- Segmentation done by user or traffic type, not geographically.



THE UNIVERSITY OF BRITISH COLUMBIA

End user experience

- “Ease of use” and “zero cost” are the two prime goals.
- Campus-wide coverage. Wireless only becomes truly useful when it’s available everywhere.
- Needs to work with any IEEE 802.11 devices with no help desk support. Calls cost money (for everyone).
- All Internet use is authenticated.
- Faculty, Students and Staff self create campus wide accounts used for many services including wireless.
- Faculty and Staff can “sponsor” guests and create accounts for them.
- Legacy devices IEEE 802.11b are here to stay; we’ll always have insecure portions to support these devices.
- But ... also need to be able to support the latest standards (especially in regards to security).



THE UNIVERSITY OF BRITISH COLUMBIA

Wireless Network Use

- email, calendaring, messaging.
- Online voting, score keeping.
- Instrumentation.
- Wireless labs anywhere, anytime.


Futures

- Utilities – cost savings in operations.
- Voice over wireless (campus wide Wi-Fi cordless phones).
- Wireless photocopier/printers.
- We've seen nothing yet, we have to be able to support any applications.



THE UNIVERSITY OF BRITISH COLUMBIA


Usage - Monthly Unique Users





THE UNIVERSITY OF BRITISH COLUMBIA

Usage - Daily Unique Users



November 11th 2003

www.wireless.ubc.ca



THE UNIVERSITY OF BRITISH COLUMBIA

Challenges

- Delays with standards-based wireless security; had expected it in 2002.
- Setting user and stakeholder expectations.
- Network fragility (although 99.6% reliable in almost two years of operation).
- Changing old models and legacy thinking.
- Vendors who aren't prepared for EDU environments.
- "Technology surfing" & changing landscape.
- Fuzzy standards with optional and incompatible technology: 802.11 FH versus DS, 802.3af options, 802.11e options.
- RF is analog, networking is digital. RF is a whole new world where 50% signal is considered good and discards are tolerated.
- Lack of true "virtual AP" capabilities.



THE UNIVERSITY OF BRITISH COLUMBIA

Key Success Factors

- User-centric service.
- Planning and research.
- Shared vision of making UBC a top University.
- UNP Project management framework.
- Dedicated wireless project team.
- A strong senior leadership and sponsorship.
- A strong online communications strategy.



THE UNIVERSITY OF BRITISH COLUMBIA

Network Design

- End-user experience drove the network design.
- Isolation of insecure wireless network using logical and physical separation.
- Segmenting by user & authentication type, not by geography.
- Segmentation using campus-wide VLANs (IEEE 802.1Q).
- Public address instead of NAT (RFC 1918) addresses; easier to track abuse.
- Broadcast management using filtering and rate limiting.
- Better broadcast is needed for Ethernet. IEEE 802.3ah will need to address large broadcast domains. Cable and DSL has already done this.
- Except for small environments, difficult to see how to justify highly proprietary & specialized cores (standards as the base technology is critical).





THE UNIVERSITY OF BRITISH COLUMBIA

Network Equipment

- 1300 Access Points – Cisco AP1200 (802.11b radio to be upgraded to 802.11g pending issue resolution).
- 200 distributions switches – Cisco 3550PWR (power over Ethernet) connected to Gigabit Ethernet (LX/SX).
- 4 core “carrier class” gigabit switches – Cisco 4507R (with redundant CPU and power).
- Web authentication servers (redundant): Colubris CN3500.
- VPN servers (to be redundant): Contivity 2700.
- Redundant Firewalls and IDS servers: Cisco PIX (& various).
- Redundant RADIUS servers: Radiator connected to six LDAP servers and then to Oracle.



THE UNIVERSITY OF BRITISH COLUMBIA

Security – still the #1 issue

- Risk Management
- Web-based authentication
- VPN Authentication and Encryption
- IEEE 802.1x authentication and encryption
- Physical security
- When solved, IEEE should co-host conference with Blackhat.com



THE UNIVERSITY OF BRITISH COLUMBIA

Web-based authentication

- Simple, users bring up browsers: automatic redirection.
- Users get started on their own: greater satisfaction and cheaper to operate. Similar to Hotspot models.
- Secure (SSL/HTTPS) authentication.
- Pass-through (unauthenticated) access to www.wireless.ubc.ca & www.library.ubc.ca
- Status/session windows provides user feedback: login ID, time and bandwidth consumption. Helps prevent abuse.
- Because HTTPS is not a stateful protocol, ARPs for duplicates and state.
- Works with any wireless client although most PDAs don't support popups for status window.



THE UNIVERSITY OF BRITISH COLUMBIA

VPN – Authentication and Encryption

- Optional but highly recommended and included free as part of the service.
- PPTPv2 support for simplicity
 - Added MSChapV2 support on LDAP
 - Many PDAs have PPTPv2 support
- IPsec for security
- Still too vulnerable to “man in the middle” at Layer 2 via ARP attacks (seen ettercap?) and other attacks.
- Can provide virtual departmental VPN services using ID followed by dot department. user.department
- VPN is not a very good technology for wireless, can’t handle fundamental wireless unreliability that well.
- Managing risks
 - How unsafe is PPTPv2?
 - How safe is IPsec in various implementations?



THE UNIVERSITY OF BRITISH COLUMBIA

IEEE 802.1x – the strategic choice for wireless authentication and encryption

- Not all EDUs are optimistic about IEEE 802.1x
- 802.1x is *still* not really deployable on a large scale without considerable pain (and costs).
- To be a success, needs to be compatible with shipping laptops and PDAs.
- WPA is a good start but Wi-Fi Alliance has no user advocacy group. They need to focus on delivering what users want not on vendor differentiation.
- We need a neutral Interoperability certification body.
- By doing login at AP, allows dynamic VLANs (equipment should not have a limited number of VLANs).
- Too many EAP variations (and increasing all the time!)



THE UNIVERSITY OF BRITISH COLUMBIA

EAP

- Pronounced “Eeeeeeeeaapppp” Definition: What implementers say to themselves when they look at the implementation issues, uncertainties and variables.
- Although implementers can control network and authentication backends, they can’t control clients.
- We need strong standards and good deployment guides for this important part of the puzzle
- The number of hours collectively wasted by implementers on EAP is a crime.
- Too many types:
 - EAP-TLS – it’s broken and should be easier to deploy
 - EAP-Cisco (LEAP) – also broken and proprietary with no support from Microsoft.
 - EAP-PEAP (is there a “standard” yet?)
 - EAP-TTLS (no Microsoft support and the permutations are multiplied)
 - EAP-SIM



THE UNIVERSITY OF BRITISH COLUMBIA

... EAP – what we really need

- Best way to do wireless authentication: Distributing limited/throw-away certificates via secure Web downloads. These could be checked across domains. My certificates could be for `martell.itservices.ubc.ca` or `martell.ca` for example.
- Would allow large wireless network operators to trust other domains. `ubc.ca` would setup trust relationships with other EDUs and with commercial Hotspot providers.
- Certificates and certificate distribution needs to be inexpensive to be ubiquitous across different platforms.
- By limiting the number of time a userID and password is used (to infrequent management of certificates), limit exposure of ID/password theft.



THE UNIVERSITY OF BRITISH COLUMBIA

Filtering

- Rogue DHCP
- SNMP filtering
- Microsoft Networking (NBT, RPC)
- Can turn on PSPF on APs and Protected Port/PVLAN on switches
- Might dramatically increase filtering when IEEE 802.1x (WPA/IEEE 802.11i) becomes deployable and/or if abuse increases.
- Exercise in risk management.



THE UNIVERSITY OF BRITISH COLUMBIA

Physical Security

- APs are the only device out of the wiring closet in typical enterprise installations
- Forcing the AP in the closet isn't ideal because of antenna cable loss and the fact that future cells might get smaller.
- Good enclosures are hard to source, most commercial ones are metal (not RF friendly).
- APs need to be able to authenticate to the switches (using IEEE 802.1x). If APs are unplugged the port is disconnected and left off. This needs to work on IEEE 802.1Q trunk ports.



THE UNIVERSITY OF BRITISH COLUMBIA

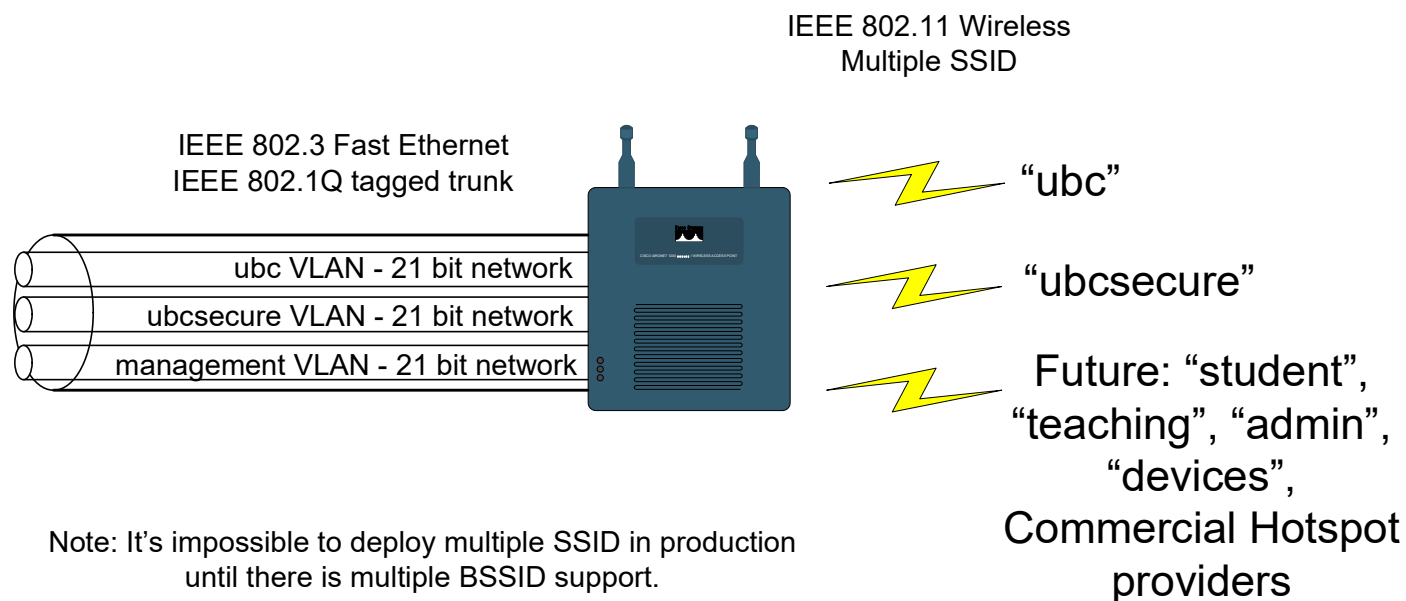
VLANs & Virtual AP support

- IEEE 802.1Q (VLAN) is a great technology.
- Currently can map multiple SSID to a single BSSID (not good enough and almost useless because of single BSSID limitations).
- Currently have two SSIDs mapped to VLANs but expect to grow to many more.
- True Virtual AP capabilities need the multiple BSSID support.
- Provides semi “out of band” management by having a higher priority protected management VLAN for all wireless devices.
- The need for dynamic VLANs. Ideal would be to have single VLAN per user and users could form groups by themselves.
- In EDU environment, we will broadcast three base wireless networks “student”, “education” and “admin”. In corporate environment, there will be a need to have a “visitor” open (but authenticated) network. VLANs support on APs is a requirement for enterprise class APs.
- Will likely keep existing “ubc” broadcast network as well as “ubcsecure” 802.1x protected network (WPA/802.11i)



THE UNIVERSITY OF BRITISH COLUMBIA

VLANs & Virtual AP support





THE UNIVERSITY OF BRITISH COLUMBIA

Authentication, Authorization and Accounting - RADIUS

- At the heart of the wireless network.
- Provides AAA services for Web login, VPN and 802.1x.
- Goes against LDAP (high availability configuration).
- Accounting info goes in enterprise SQL databases.
- Track user ID, machine/Mac, IP, bytes, time (critical to get hogs and other abusers).



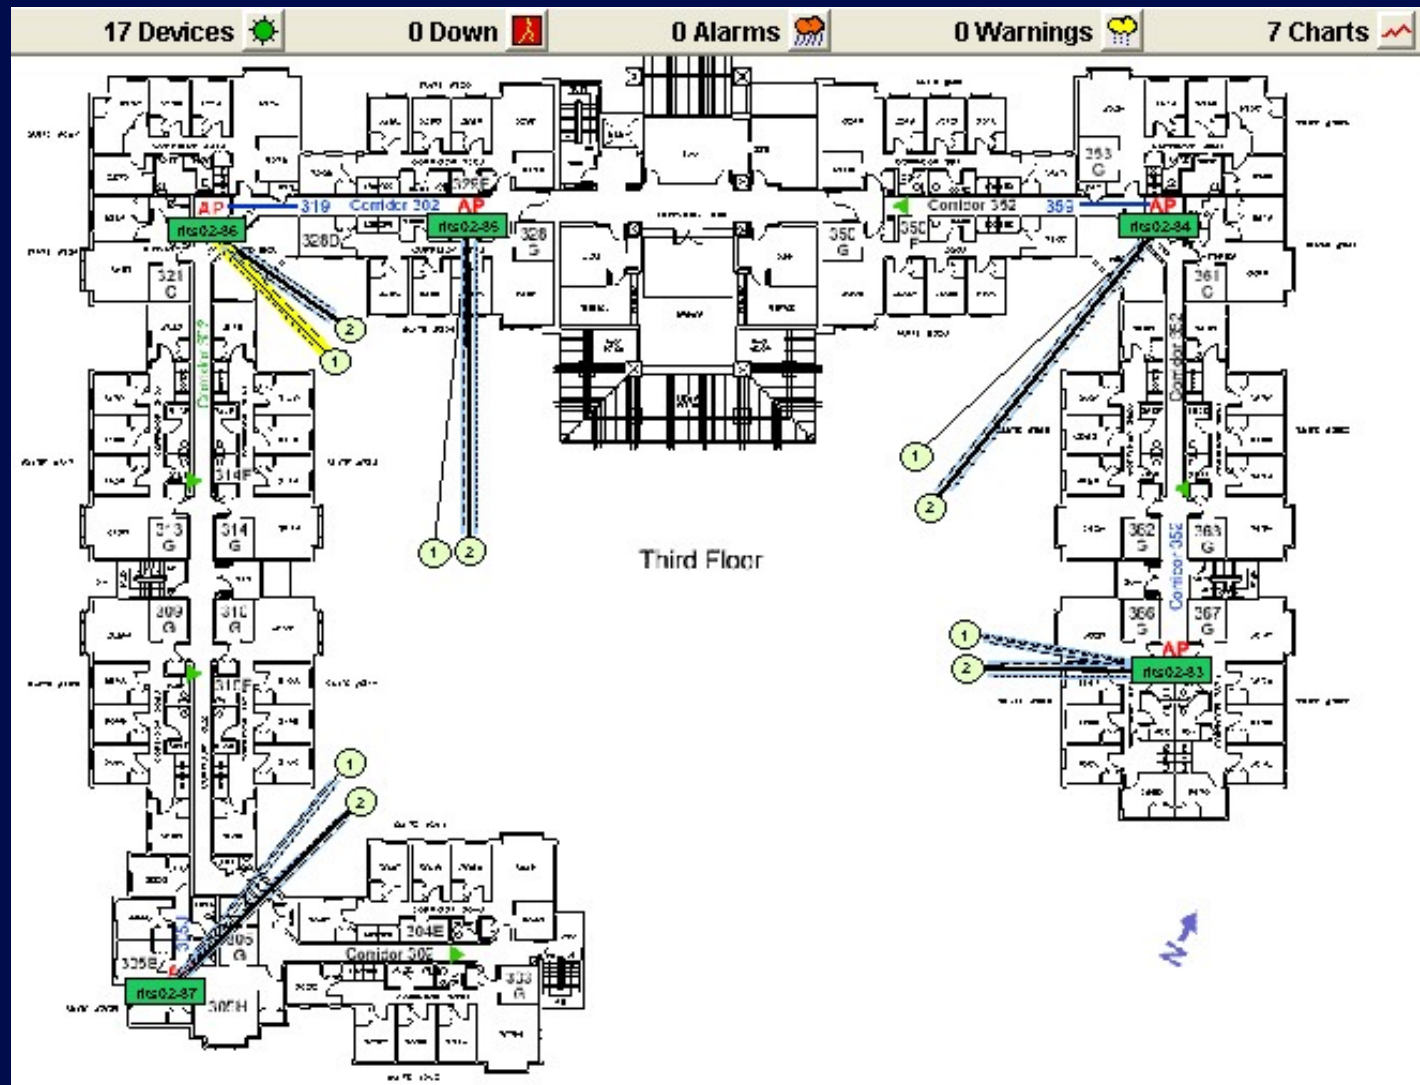
THE UNIVERSITY OF BRITISH COLUMBIA

Management

- Devices are on a segmented VLAN not accessible from user or wireless networks.
- Vendors tools aren't there yet for large networks but we have an "off the shelf" network.
- Lightweight tools versus expensive, complex and inflexible heavyweights.
- All network devices also documented in databases.
- Extensive SNMP based management via scripts and Intermapper
- WLSE – Cisco's AP management tool used to assist in RF data collection. Needs to have programmable interface. WDS needs to be ported to core switches.
- Need physical security of AP or AP acting as 802.1x client to switches.



THE UNIVERSITY OF BRITISH COLUMBIA

November 11th 2003

www.wireless.ubc.ca

29



THE UNIVERSITY OF BRITISH COLUMBIA


Total WAP in database: 1217

DEVICE_ID	DEVICE_NAME	LOC_BLDG_ID	LOC_ROOM_NO	BASE_IP_ADDR	BASE_MAC_ADDR	POWER_TYPE_CDE	RADIO_ONE_TYPE	RADIO_I
99	CISR02-99	165	8	142.103.28.169	000A8A77D2A6	POE-CSCO	802.11B	000750D
98	CISR02-98	165	14	142.103.28.168	000A8A77D2A3	POE-CSCO	802.11B	000750D
97	CISR02-97	165	39	142.103.28.167	000A8A77D29A	POE-CSCO	802.11B	000750D
96	CISR02-96	165	51	142.103.28.166	000A8A8ACC20	POE-CSCO	802.11B	000750D
95	CISR02-95	165	67	142.103.28.165	000A8A77D4A2	POE-CSCO	802.11B	000750D
94	CISR02-94	165	92	142.103.28.164	000A8A8AC992	POE-CSCO	802.11B	000750D
93	CISR02-93	165	108	142.103.28.163	000A8A8ACB3B	POE-CSCO	802.11B	000750D
92	CISR02-92	165	144	142.103.28.162	000A8A8ACB24	POE-CSCO	802.11B	000750D
90	CISR02-90	165	180	142.103.28.160	000A8A8ACB2F	POE-CSCO	802.11B	000750D
88	CISR02-88	165	240	142.103.28.158	000A8A8ACC33	POE-CSCO	802.11B	000750D
86	CISR02-86	165	288	142.103.28.156	000A8A8AC922	POE-CSCO	802.11B	000750D
84	CISR02-84	165	342	142.103.28.154	000A8A8ACA1C	POE-CSCO	802.11B	000750D
83	CISR02-83	165	350	142.103.28.153	000A8A8ACA28	POE-CSCO	802.11B	000750D
82	CISR02-82	165	380	142.103.28.152	000A8A8ACA27	POE-CSCO	802.11B	000750D
91	CISR02-91	165	166	142.103.28.161	000A8A8ACB23	POE-CSCO	802.11B	000750D
89	CISR02-89	165	219	142.103.28.159	000A8A8ACC23	POE-CSCO	802.11B	000750D
87	CISR02-87	165	262	142.103.28.157	000A8A8AC931	POE-CSCO	802.11B	000750D
85	CISR02-85	165	308	142.103.28.155	000A8A8ACA37	POE-CSCO	802.11B	000750D
99	BRCK02-99	112-1	2036	142.103.30.247	000C309D4633	POE-CSCO	802.11B	000785B
98	BRCK02-98	112-1	2016	142.103.30.246	000C309D4786	POE-CSCO	802.11B	000C306
97	BRCK02-97	112-1	2073	142.103.30.245	000C30DA9DAB	POE-CSCO	802.11B	000C306
96	BRCK02-96	112-1	1040	142.103.30.244	000C309D4687	POE-CSCO	802.11B	000785B
95	BRCK02-95	112-1	1073	142.103.30.243	000C309D46D3	POE-CSCO	802.11B	000C306
94	BRCK02-94	112-1	1020	142.103.30.242	000C309D4790	POE-CSCO	802.11B	000C306
93	BRCK02-93	112-1	1002	142.103.30.241	000C309D4718	POE-CSCO	802.11B	000C306



THE UNIVERSITY OF BRITISH COLUMBIA

Managing concurrent use



November 11th 2003

www.wireless.ubc.ca

31



THE UNIVERSITY OF BRITISH COLUMBIA

Supporting Research

- Massive test bed.
- The need to balance operations with research.
- Developed of Visual Mapping tool for recording survey information.
- Automatic channel and power assignment technology under development (which we hope to “plug-in” to WLSE)
- Other propagation studies underway:
 - diversity antennas?
 - overlapping channels?
 - does a spectrum need to be managed to be reliable?



THE UNIVERSITY OF BRITISH COLUMBIA

Impact of newer standards

- IEEE 802.11e? - QOS will be done on a per-VLAN basis. Telephony, transaction and management wireless vlans need highest priority. QOS threat is not from friendly or well behaved RF.
- Fast Roaming – an absolute minimum to support and scale mobility. Need a good solution both at Layer 2 first and then at Layer 3 (and between other technology like 802.20)
- 802.11k - Radio resources: should also work to provide client assisted information and detect rogues & interference. Vendor implementations will likely lead the way in the short term.
- Ideally, equipment should be able to run (with regulatory unlock code) on other spectrum around 2.4 GHz



THE UNIVERSITY OF BRITISH COLUMBIA

Future - Spectrum

- Success of low cost, unlicensed spectrum is clear... but reliability and spectrum congestion is an issue.
- Expanding the spectrum: the FCC & [Industry Canada](#)
 - 2001 Speech from the Throne: "...making broadband access widely available to citizens, businesses, public institutions and to all communities by 2005."
- 2nd generation high speed wireless technology should provide reliable and cost effective networking using:
 - commodity products
 - low power regulations (smaller "markets")
 - inexpensive "managed" spectrum
 - municipal and campus licenses
- IEEE 802.20 – future high speed mobile broadband?



THE UNIVERSITY OF BRITISH COLUMBIA

Spectrum auctions are not the solution

- Current Industry Canada logic
 - “Auctions offer a number of advantages over the other options that are available to governments to assign access to the radio spectrum such as: their ability to promote economically efficient use of spectrum; their openness and objectivity as an assignment mechanism; their procedural efficiency; and their ability to return appropriate compensation to Canadian taxpayers for the use of a public resource.”
 - “The Government’s objective in conducting auctions is not to raise revenue, but rather to award licences fairly, efficiently and effectively so as to ensure that the Canadian public derives the maximum possible benefit from the spectrum resource.”
 - “Auction bids thus depend on consumer prices; consumer prices do not depend on auction bids.” [Reference](#)
- This logic doesn’t work when if consumer price goal is “zero cost”.



THE UNIVERSITY OF BRITISH COLUMBIA

Questions

- Information on UBC Network
www.wireless.ubc.ca
- Email contact:
 - Work: jonm.martell@ubc.ca
 - Personal: jonm@martell.ca

EXHIBIT D

FEB 7 2007

Dictionary of Computer and Internet Terms

Ninth Edition


Douglas A. Downing, Ph.D.
School of Business and Economics
Seattle Pacific University

Michael A. Covington, Ph.D.
Artificial Intelligence Center
The University of Georgia

Melody Mauldin Covington
Covington Innovations
Athens, Georgia

With the assistance of
Catherine Anne Covington

FOR REFERENCE ONLY
Do Not Take From This Room



ABOUT THE AUTHORS

Douglas Downing teaches economics and quantitative methods and is undergraduate program director for the School of Business and Economics at Seattle Pacific University. He is the author of several books in both Barron's Easy Way and Business Review series. He is also the author of *Java Programming the Easy Way* and *Dictionary of Mathematics Terms*, published by Barron's Educational Series, Inc. He holds the Ph.D. degree in economics from Yale University.

Michael Covington is Associate Director of the Artificial Intelligence Center at the University of Georgia. He is the author of several books and over 250 magazine articles. He holds the Ph.D. degree in linguistics from Yale University.

Melody Mauldin Covington is a graphic designer living in Athens, Georgia. She is the author of *Dictionary of Desktop Publishing* (published by Barron's).

Catherine Anne Covington is a student and web designer in Athens, Georgia.

© Copyright 2006, 2003, 2000, 1998, 1996, 1995, 1992, 1989, and 1986
by Barron's Educational Series, Inc.

All rights reserved.

No part of this book may be reproduced in any form, by photostat, microfilm, xerography, or any other means, or incorporated into any information retrieval system, electronic or mechanical, without the written permission of the copyright owner.

All inquiries should be addressed to:

Barron's Educational Series, Inc.

250 Wireless Boulevard

Hauppauge, New York 11788

<http://www.barronseduc.com>

Library of Congress Catalog Card No. 2005052175

ISBN-13: 978-0-7641-3417-3

ISBN-10: 0-7641-3417-5

Library of Congress Cataloging-in-Publication Data

Downing, Douglas.

Dictionary of computer and Internet terms / Douglas A. Downing, Michael A. Covington, Melody Mauldin Covington.—9th ed.

p. cm.

ISBN-13: 978-0-7641-3417-3

ISBN-10: 0-7641-3417-5

1. Computers—Dictionaries. 2. Internet—Dictionaries. I. Covington, Michael A., 1957– II. Covington, Melody Mauldin. III. Title.

QA76.15.D667 2006

004'.003—dc22

2005052175

PRINTED IN THE UNITED STATES OF AMERICA

987654321

D

DAC, D/A converter *see* DIGITAL-TO-ANALOG CONVERTER.

daemon (under UNIX) a program that runs continuously in the background, or is activated by a particular event. The word *daemon* is Greek for “spirit” or “soul.”

dagger the character †, sometimes used to mark footnotes. *See also* FOOTNOTE. Also called an OBELISK or LONG CROSS.

daisy-chain to connect devices together in sequence with cables. For example, if four devices A, B, C, and D are daisy-chained, there will be a cable from A to B, a cable from B to C, and a cable from C to D.

daisywheel printer a printer that uses a rotating plastic wheel as a type element. Daisywheel printers were often used with microcomputers in the early 1980s. They printed high-quality text, but they were relatively slow and could not print graphics.

dash (—) a punctuation mark similar to a hyphen, but longer. On a typewriter, a dash is typed as two hyphens.

Proportional-pitch type often includes one or more kinds of dashes, such as an em dash (—), which is as wide as the height of the font, and an en dash (–), which is two-thirds as wide as the em dash. Normally, the em dash joins sentences and the en dash joins numbers (as in “1995–98”).

data information. The word was originally the plural of *datum*, which means “a single fact,” but it is now often used as a collective singular. Data processing is the act of using data for making calculations or decisions. *Usage note:* This usage came and went.

database a collection of data stored on a computer storage medium, such as a disk, that can be used for more than one purpose. For example, a firm that maintains a database containing information on its employees will be able to use the same data for payroll, personnel, and other purposes. *See* DATABASE MANAGEMENT.

database management the task of storing data in a database and retrieving information from that data. There are three aspects of database management: entering data, modifying or updating data, and presenting output reports. Many mainframe computers are used by businesses for database management purposes. Several software packages are available for database management on microcomputers, such as dBASE and Microsoft Access, and some data management capabilities are provided with spreadsheets such as Lotus 1-2-3 and Excel. Some examples of database applications include maintaining employee lists and preparing payrolls; maintaining parts order lists and keeping

McGRAW-HILL DICTIONARY OF SCIENTIFIC AND TECHNICAL TERMS

**Sixth
Edition**

McGraw-Hill

New York Chicago San Francisco
Lisbon London Madrid Mexico City
Milan New Delhi San Juan Seoul Singapore Sydney Toronto

On the cover: Representation of a fullerene molecule with a noble gas atom trapped inside. At the Permian-Triassic sedimentary boundary the noble gases helium and argon have been found trapped inside fullerenes. They exhibit isotope ratios quite similar to those found in meteorites, suggesting that a fireball meteorite or asteroid exploded when it hit the Earth, causing major changes in the environment. (Image copyright © Dr. Luann Becker. Reproduced with permission.)

Over the six editions of the Dictionary, material has been drawn from the following references: G. M. Garrity et al., *Taxonomic Outline of the Prokaryotes*, Release 2, Springer-Verlag, January 2002; D. W. Linzey, *Vertebrate Biology*, McGraw-Hill, 2001; J. A. Pechenik, *Biology of the Invertebrates*, 4th ed., McGraw-Hill, 2000; U.S. Air Force *Glossary of Standardized Terms*, AF Manual 11-1, vol. 1, 1972; F. Casey, ed., *Compilation of Terms in Information Sciences Technology*, Federal Council for Science and Technology, 1970; *Communications-Electronics Terminology*, AF Manual 11-1, vol. 3, 1970; P. W. Thrush, comp. and ed., *A Dictionary of Mining, Mineral, and Related Terms*, Bureau of Mines, 1968; *A DOD Glossary of Mapping, Charting and Geodetic Terms*, Department of Defense, 1967; J. M. Gilliland, *Solar-Terrestrial Physics: A Glossary of Terms and Abbreviations*, Royal Aircraft Establishment Technical Report 67158, 1967; W. H. Allen, ed., *Dictionary of Technical Terms for Aerospace Use*, National Aeronautics and Space Administration, 1965; *Glossary of Stirling Terminology*, Office of Aerospace Research, U.S. Air Force, 1963; *Naval Dictionary of Electronic, Technical, and Imperative Terms*, Bureau of Naval Personnel, 1962; R. E. Huschke, *Glossary of Meteorology*, American Meteorological Society, 1959; *ADP Glossary*, Department of the Navy, NAVSO P-3097; *Glossary of Air Traffic Control Terms*, Federal Aviation Agency; *A Glossary of Range Terminology*, White Sands Missile Range, New Mexico, National Bureau of Standards, AD 467-424; *Nuclear Terms: A Glossary*, 2d ed., Atomic Energy Commission.

McGRAW-HILL DICTIONARY OF SCIENTIFIC AND TECHNICAL TERMS, Sixth Edition

Copyright © 2003, 1994, 1989, 1984, 1978, 1976, 1974 by The McGraw-Hill Companies, Inc. All rights reserved. Printed in the United States of America. Except as permitted under the United States Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written permission of the publisher.

1 2 3 4 5 6 7 8 9 0 DOW/DOW 0 8 7 6 5 4 3 2

ISBN 0-07-042313-X

Library of Congress Cataloging-in-Publication Data

McGraw-Hill dictionary of scientific and technical terms--6th ed.

p. cm.

ISBN 0-07-042313-X (alk. paper)

1. Science--Dictionaries. 2. Technology--Dictionaries. I. Title: Dictionary of scientific and technical terms.

Q123.M15 2002
503--dc21

2002026436

anaconda [VERT ZOO] *Eunectes murinus*. The largest living snake, an arboreal-aquatic member of the boa family (Boidae). { 'an-ə-'kän-də }

anacoustic zone [GEOPHYS] The zone of silence in space, starting at about 100 miles (160 kilometers) altitude, where the distance between air molecules is greater than the wavelength of sound, and sound waves can no longer be propagated. { 'an-ə-'kū-stik, zōn }

Anactinotrichosel [INV ZOO] A group name for three closely related suborders of mites and ticks: Onychopalpida, Mesostigmata, and Ixodidae. { 'an-ə-'tō-nə-'kīt-ən-'dō-sē }

Anacystis [BOT] A genus of blue-green algae in the class Cyanophyceae. { 'an-ə-'sīs-təs }

anadromous [VERT ZOO] Said of a fish, such as the salmon and shad, that ascends fresh-water streams from the sea to spawn. { 'ə-'nā-drə-məs }

Anadyomenaceae [BOT] A family of green marine algae in the order Siphonocladales characterized by the expanded blades of the thallus. { 'ə-'nā-dy-ə-'men-'ās-ē-ē }

anaerobe [BIOL] An organism that does not require air or free oxygen to maintain its life processes. { 'an-ə-'rōb }

anaerobic adhesive [MATER] A single-component adhesive that hardens rapidly to form a strong bond between surfaces from which air is excluded. { 'an-ə-'rōb-ik əd'hē-ziv }

anaerobic bacteria [MICROBIO] Any bacteria that can survive in the partial or complete absence of air; two types are facultative and obligate. { 'an-ə-'rōb-ik, 'bāk-tīr-ē-ə }

anaerobic condition [BIOL] The absence of oxygen, preventing normal life for organisms that depend on oxygen. { 'an-ə-'rōb-ik kən'dish-ən }

anaerobic glycolysis [BIOCHEM] A metabolic pathway in plants by which, in the absence of oxygen, hexose is broken down to lactic acid and ethanol with some adenosinetriphosphate synthesis. { 'an-ə-'rōb-ik glī-'kāl-ə-'səs }

anaerobic petri dish [MICROBIO] A glass laboratory dish for plate cultures of anaerobic bacteria; a thioglycollate agar medium and restricted air space give proper conditions. { 'an-ə-'rōb-ik 'pē-trē, 'dīsh }

anaerobic process [SCI TECH] A process from which air or oxygen not in chemical combination is excluded. { 'an-ə-'rōb-ik 'prās-əs }

anaerobic sediment [GEOL] A highly organic sediment formed in the absence or near absence of oxygen in water that is rich in hydrogen sulfide. { 'an-ə-'rōb-ik 'sed-ə-'mənt }

anaerobiosis [BIOL] A mode of life carried on in the absence of molecular oxygen. { 'an-ə-'rōb-ē-'səs }

anaerophyte [ECOL] A plant that does not need free oxygen for respiration. { 'ə-'ner-ə-'fīt }

anafont [METEOROL] A front at which the warm air is ascending the frontal surface up to high altitudes. { 'an-ə-'frənt }

anagen effluvium [MED] Acute hair loss that usually follows chemotherapy or radiotherapy. { 'an-ə-'jən ə'flū-'vē-əm }

anaglyph [GRAPHICS] 1. A stereogram in which the two views are printed or projected superimposed in complementary colors, usually red and blue; by viewing through filter spectacles of corresponding complementary colors, a stereoscopic image is formed. 2. A surface worked in low relief. { 'an-ə-'glīf }

anagryne [ORG CHEM] $C_{15}H_{20}N_2O$ A toxic alkaloid found in several species of *Lupinus* in the western United States; acute poisoning produces nervousness, depression, loss of muscular control, convulsions, and coma. { 'an-ə-'jī-rēn }

anaknesis [BIOCHEM] A process in living organisms by which energy-rich molecules, such as adenosine triphosphate, are formed. { 'an-ə-'kə-nēs-əs }

anal [ANAT] Relating to or located near the anus. { 'ān-əl }

analbite [MINERAL] A triclinic albite which is not stable and becomes monoclinic at about 700°C. { 'ə-'nāl,bīt }

albuminemia [MED] A disorder transmitted as an autosomal recessive, characterized by drastic reduction or absence of serum albumin. { 'an,əl,byū-'mā-nēm-ē-ə }

anal character [PSYCH] A personality type that exhibits excessive orderliness, miserliness, and obstinacy. { 'ān-əl 'kār-ik-tər }

analclime [MINERAL] $NaAlSi_2O_6 \cdot H_2O$ A white or slightly colored isometric zeolite found in diabase and in alkali-rich basalts. Also known as analcite. { 'ə-'nāl,sēm }

analclimite [PETR] An extrusive or hypabyssal rock that consists primarily of pyroxene and analclime. { 'ə-'nāl-sə,mīt }

analclimzation [GEOL] The replacement in igneous rock of feldspars or feldspathoids by analclime. { 'ə-'nāl-sə-'mā-'zā-shən }

analclite See analclime. { 'ə-'nāl,sīt }

analemma [ASTRON] A figure-eight-shaped diagram on a globe showing the declination of the sun throughout the year and also the equation of time. [CIV ENG] Any raised construction which serves as a support or rest. { 'an-ə-'lem-ə }

analeptic [PHARM] Any drug used to restore respiration and a wakeful state. { 'an-ə-'lep-tik }

anal fin [VERT ZOO] An unpaired fin located medially on the posterior ventral part of the fish body. { 'ān-əl, 'fīn }

analgesia [PHYSIO] Insensibility to pain with no loss of consciousness. { 'an-əl-'jēzh-ə }

analgesic [PHARM] Any drug, such as salicylates, morphine, or opiates, used primarily for the relief of pain. { 'an-əl-'jēz-ik }

anal gland [INV ZOO] A gland in certain mollusks that secretes a purple substance. [VERT ZOO] A gland located near the anus or opening into the rectum in many vertebrates. { 'ān-əl, 'glānd }

anallagmatic curve [MATH] A curve that is its own inverse curve with respect to some circle. { 'ə-'nāl-'ig-mad-ik 'kərv }

anallobaric center See pressure-rise center. { 'ə-'nāl-ə-'bār-ik 'sen-tər }

analog [CHEM] A compound whose structure is similar to that of another compound but whose composition differs by one element. [FOOD ENG] A meat-substitute food manufactured from vegetable ingredients, such as soybeans. [ELECTR] 1. A physical variable which remains similar to another variable insofar as the proportional relationships are the same over some specified range; for example, a temperature may be represented by a voltage which is its analog. 2. Pertaining to devices, data, circuits, or systems that operate with variables which are represented by continuously measured voltages or other quantities. [METEOROL] A past large-scale synoptic weather pattern which resembles a given (usually current) situation in its essential characteristics. { 'an-əl, 'æg }

analog adder [ELECTR] A device with one output voltage which is a weighted sum of two input voltages. { 'an-əl, 'æg 'ad-ər }

analog channel [ELECTR] A channel on which the information transmitted can have any value between the channel limits, such as a voice channel. { 'an-əl, 'æg 'chan-əl }

analog communications [COMMUN] System of telecommunications employing a nominally continuous electric signal that varies in frequency, amplitude, or other characteristic, in some direct correlation to nonelectrical information (sound, light, and so on) impressed on a transducer. { 'an-əl, 'æg kə'myū-nə-'kā-shənz }

analog comparator [ELECTR] 1. A comparator that checks digital values to determine whether they are within predetermined upper and lower limits. 2. A comparator that produces high and low digital output signals when the sum of two analog voltages is positive and negative, respectively. { 'an-əl, 'æg kəm'pār-əd-ər }

analog computer [COMPUT SCI] A computer in which quantities are represented by physical variables; problem parameters are translated into equivalent mechanical or electrical circuits as an analog for the physical phenomenon being investigated. { 'an-əl, 'æg kəm'pyūd-ər }

analog data [COMPUT SCI] Data represented in a continuous form, as contrasted with digital data having discrete values. { 'an-əl, 'æg 'dād-ə }

analog-digital computer See hybrid computer. { 'an-əl, 'æg 'dij-ə-təl kəm'pyūd-ər }

analog indicator [ELECTR] A device in which the result of a measurement is indicated by a pointer deflection or other visual quantity. { 'an-əl, 'æg 'in-də-'kād-ər }

analog monitor [ELECTR] A display unit that accepts only analog signals, which must be converted from digital signals by the computer's video display board. { 'an-əl, 'æg 'mān-əd-ər }

analog multiplexer [ELECTR] A multiplexer that provides switching of analog input signals to allow use of a common analog-to-digital converter. { 'an-əl, 'æg 'məl-tə-'plek-sər }


analog multiplier [ELECTR] A device that accepts two or more inputs in analog form and then produces an output proportional to the product of the input quantities. { 'an-əl, 'æg 'məl-tə-'plī-ər }

ANADYOMENACEAE



Anadyomene, a genus in Anadyomenaceae, with expanded blades.

ANAEROBIC PETRI DISH



Brewer anaerobic petri dish. (Courtesy BioQuest, Division of Becton, Dickinson and Co.)

DAT See digital audio tape.

data [COMPUT SCI] 1. General term for numbers, letters, symbols, and analog quantities that serve as input for computer processing. 2. Any representations of characters or analog quantities to which meaning, if not information, may be assigned. [SCI TECH] Numerical or qualitative values derived from scientific experiments. { 'dad·ə, 'dād·ə, or 'dād·ə }

data acquisition [COMMUN] The phase of data handling that begins with the sensing of variables and ends with a magnetic recording or other record of raw data; may include a complete radio telemetering link. { 'dad·ə, 'ak·wə,zish·ən }

data acquisition computer [COMPUT SCI] A computer that is used to acquire and analyze data generated by instruments. { 'dad·ə, 'ak·wə,zish·ən kəm'pyüt·ər }

data aggregate [COMPUT SCI] The set of data items within a record. { 'dad·ə, 'ag·rə·gət }

data analysis [COMPUT SCI] The evaluation of digital data. { 'dad·ə, 'nal·ə·səs }

data attribute [COMPUT SCI] A characteristic of a block of data, such as the type of representation used or the length in characters. { 'dad·ə, 'a·trə'byüt }

data automation [COMPUT SCI] The use of electronic, electromechanical, or mechanical equipment and associated techniques to automatically record, communicate, and process data and to present the resultant information. { 'dad·ə, 'əd·ə'mā·shən }

data bank [COMPUT SCI] A complete collection of information such as contained in automated files, a library, or a set of computer disks. { 'dad·ə, 'bāŋk }

database [COMPUT SCI] A nonredundant collection of interrelated data items that can be shared and used by several different subsystems. { 'dad·ə, 'bās }

database/data communication [COMPUT SCI] An advanced software product that combines a database management system with data communications procedures. Abbreviated DB/DC. { 'dad·ə, 'bās 'dad·ə kə,myū·nə'kā·shən }

database machine [COMPUT SCI] A computer that handles the storage and retrieval of data into and out of a database. { 'dad·ə, 'bās mə,shēn }

database management system [COMPUT SCI] A special data processing system, or part of a data processing system, which aids in the storage, manipulation, reporting, management, and control of data. Abbreviated DBMS. { 'dad·ə, 'bās 'man·ij·mənt, 'sis·təm }

database server [COMPUT SCI] An independently functioning computer in a local-area network that holds and manages the database. { 'dad·ə, 'bās, 'sər·vər }

data break [COMPUT SCI] A facility which permits input/output transfers to occur without disturbing program execution in a computer. { 'dad·ə, 'brāk }

data buffering [COMPUT SCI] The temporary collection and storage of data awaiting further processing in physical storage devices, allowing a computer and its peripheral devices to operate at different speeds. { 'dad·ə, 'bʌf·ə·rɪŋ }

data bus [ELECTR] An internal channel that carries data between a computer's central processing unit and its random-access memory. { 'dad·ə, 'bʌs }

data capture [COMPUT SCI] The acquisition of data to be entered into a computer. { 'dad·ə, 'kæp·tʃər }

data carrier [COMPUT SCI] A medium on which data can be recorded, and which is usually easily transportable, such as cards, tape, paper, or disks. { 'dad·ə, 'kɑ·r·ē·ər }

data carrier storage [COMPUT SCI] Any type of storage in which the storage medium is outside the computer, such as tape, cards, or disks, in contrast to inherent storage. { 'dad·ə, 'kɑ·r·ē·ər, 'stɔ·rɪj }

data cartridge [COMPUT SCI] A tape cartridge used for non-volatile and removable data storage in small digital systems. { 'dad·ə, 'kɑ·r·trɪj }

data cell drive [COMPUT SCI] A large-capacity storage device consisting of strips of magnetic tape which can be individually transferred to the read-write head. { 'dad·ə, 'sel, 'drɪv }

data center [COMPUT SCI] An organization established primarily to acquire, analyze, process, store, retrieve, and disseminate one or more types of data. { 'dad·ə, 'sen·tər }

data chain [COMPUT SCI] Any combination of two or more data elements, data items, data codes, and data abbreviations

in a prescribed sequence to yield meaningful information; for example, "date" consists of data elements year, month, and day. { 'dad·ə, 'tʃæn }

data chaining [COMPUT SCI] A technique used in scatter reading or scatter writing in which new storage areas are defined for use as soon as the current data transfer is completed. { 'dad·ə, 'tʃæn·ɪŋ }

data channel [COMPUT SCI] A bidirectional data path between input/output devices and the main memory of a digital computer permitting one or more input/output operations to proceed concurrently with computation. { 'dad·ə, 'tʃæn·əl }

data circuit [ELECTR] A telephone facility that allows transmission of digital data pulses with minimum distortion. { 'dad·ə, 'sər·kət }

data code [COMPUT SCI] A number, letter, character, symbol, or any combination thereof, used to represent a data item. { 'dad·ə, 'kɔd }

data collection [COMPUT SCI] The process of sending data to a central point from one or more locations. { 'dad·ə, 'kə, 'lek·shən }

data communication network [COMPUT SCI] A set of nodes, consisting of computers, terminals, or some type of communication control units in various locations, connected by links consisting of communication channels providing a data path between the nodes. { 'dad·ə, 'kə, 'myū·nə, 'kā·shən 'net, wɜrk }

data communications [COMMUN] The conveying from one location to another by electrical means of information that originates or is recorded in alphabetic, numeric, or pictorial form, or as a signal that represents a measurement; includes telemetering, telegraphy, and facsimile but not voice or television. Also known as data transmission. { 'dad·ə, 'kə, 'myū·nə'kā·shənz }

data communications processor [COMPUT SCI] A small computer used to control the flow of data between machines and terminals over communications channels. { 'dad·ə, 'kə, 'myū·nə'kā·shənz, 'prəs·es·ər }

data compression [COMPUT SCI] The technique of reducing the number of binary digits required to represent data. { 'dad·ə, 'kəm, 'preʃ·ən }

data concentrator [ELECTR] A device, such as a microprocessor, that takes data from several different teletypewriter or other slow-speed lines and feeds them to a single higher-speed line. { 'dad·ə, 'kən·sən, 'trād·ər }

data conversion [COMPUT SCI] The changing of the representation of data from one form to another, as from binary to decimal, or from one physical recording medium to another, as from card to disk. Also known as conversion. { 'dad·ə, 'kən, 'vər·zhən }

data conversion line [COMPUT SCI] The channel, electronic or manual, through which data elements are transferred between data banks. { 'dad·ə, 'kən, 'vər·zhən, 'lɪn }

data converter See converter. { 'dad·ə, 'kən, 'vərd·ər }

data definition [COMPUT SCI] The statements in a computer program that specify the physical attributes of the data to be processed, such as location and quantity of data. { 'dad·ə, 'def·ə'nɪʃ·ən }

data dependence graph [COMPUT SCI] A chart that represents a program in a data flow language, in which each node is a function and each arc carries a value. { 'dad·ə, 'dɪ, 'pen·dəns, 'graf }

data description language [COMPUT SCI] A programming language used to specify the arrangement of data items within a data base. { 'dad·ə, 'dɪ'skrɪp·shən, 'læŋ·gwɪj }

data descriptor [COMPUT SCI] A pointer indicating the memory location of a data item. { 'dad·ə, 'dɪ'skrɪp·tər }

data dictionary [COMPUT SCI] A catalog which contains the names and structures of all data types. { 'dad·ə, 'dɪk·shən, 'ner·ē }

data display [COMPUT SCI] Visual presentation of processed data by specially designed electronic or electromechanical devices through interconnection (either on- or off-line) with digital computers or component equipments; although line printers and punch cards may display data, they are not usually categorized as displays but as output equipments. { 'dad·ə, 'dɪ, 'splə }

data distribution [COMPUT SCI] Data transmission to one or more locations from a central point. { 'dad·ə, 'dɪs·trɪ, 'byū·shən }

diffusivity analysis

The quantity of heat passing normally through a unit area per unit time divided by the product of specific heat, density, and temperature gradient. Also known as thermal diffusivity; thermometric conductivity. {dif-yu'ziv-əd-ē}

diffusivity analysis [ANALY CHEM] Analysis of difficult-to-separate materials in solution by diffusion effects, using, for example, dialysis, electroanalysis, interferometry, amperometric titration, polarography, or voltammetry. {dif-yu'ziv-əd-ē ə'nal-ə'səs}

difunctional molecule [ORG CHEM] An organic structure possessing two sites that are highly reactive. {di'fəŋk-shən-əl mäl-ə-kyul}

digallic acid See tannic acid. {di'gal-ik 'as-əd}

digamma function [MATH] The derivative of the natural logarithm of the gamma function. {di'gam-ə 'fəŋk-shən}

digastric [ANAT] Of a muscle, having a fleshy part at each end and a tendinous part in the middle. {di'gas-trik}

Digenea [INV ZOO] A group of parasitic flatworms or flukes constituting a subclass or order of the class Trematoda and having two types of generations in the life cycle. {di'jē-nē-ə}

digensis [BIOL] Sexual and asexual reproduction in succession. {di'jē-n-ə'səs}

diginite [MINERAL] Cu_2S_3 A blue to black mineral consisting of an isometric copper sulfide having a variable deficiency in copper. Also known as alpha chalcocite; blue chalcocite. {di'jē-nīt}

Di George's syndrome See thymic aplasia. {də'jōrj-əz 'sin-drəm}

digested sludge [CIV ENG] Sludge or thickened mixture of sewage solids with water that has been decomposed by anaerobic bacteria. {də'jes-təd 'sləj}

digester [CHEM ENG] A vessel used to produce cellulose pulp from wood chips by cooking under pressure. [CIV ENG] A sludge-digestion tank containing a system of hot water or steam pipes for heating the sludge. {də'jes-tər}

digestion [CHEM ENG] 1. Preferential dissolving of mineral constituents in concentrations of ore. 2. Liquefaction of organic waste materials by action of microbes. 3. Separation of fabric from tires by the use of hot sodium hydroxide. 4. Removing lignin from wood in manufacture of chemical cellulose paper pulp. [CIV ENG] The process of sewage treatment by the anaerobic decomposition of organic matter. [PHYSIO] The process of converting food to an absorbable form by breaking it down to simpler chemical compounds. {də'jes-chən}

digestive efficiency [ECOL] A measure of the amount of ingested chemical energy actually absorbed by an animal. {di'jes-tiv i'fish-ən-sē}

digestive enzyme [BIOCHEM] Any enzyme that causes or aids in digestion. {də'jes-tiv 'en-zīm}

digestive gland [PHYSIO] Any structure that secretes digestive enzymes. {də'jes-tiv 'glānd}

digestive system [ANAT] A system of structures in which food substances are digested. {də'jes-tiv 'sis-təm}

digestive tract [ANAT] The alimentary canal. {də'jes-tiv 'trakt}

digger [ENG] A tool or apparatus for digging in the ground. [MIN ENG] A person who digs in the ground; usually refers to a coal miner. {di'g-ər}

digging [ENG] A sudden increase in cutting depth of a cutting tool due to an erratic change in load. {di'g-iŋ}

digging height See bank height. {di'g-iŋ 'hīt}

digging line See inhaul cable. {di'g-iŋ 'līn}

diggings [SCI TECH] 1. Excavated materials. 2. A place of excavating. {di'g-iŋz}

diglitrin [BIOCHEM] $C_{21}H_{21}O_{10}$ A flavone compound that is found in foxglove leaves. {di'g-lī-trin}

digloom [COMMUN] A wire communication system that transmits speech signals in the form of corresponding trains of pulses and transmits digital information directly from computers, radar, tape readers, teleprinters, and telemetering equipment. {di'g-lōm}

diglcon [ELECTR] An image tube in which the image produced by electrons from the photocathode is focused directly on a silicon diode array and each incoming photoelectron produces an electrical pulse that is amplified and recorded. {di'g-lōn}

digit [COMPUT SCI] In a decimal digital computer, the space reserved for storage of one digit of information. [MATH] A

character used to represent one of the nonnegative integers smaller than the base of a system of positional notation. Also known as numeric character. {di'j-ət}

digit absorbing selector [ELECTR] Dial switch arranged to set up and then fall back on the first one of two digits dialed; it then operates on the next digit dialed. {di'j-ət əb'sorb-iŋ si'lek-tər}

digital [COMPUT SCI] Pertaining to data in the form of digits. {di'j-əd-əl}

digital audio broadcasting [COMMUN] The radio broadcasting of audio signals encoded in digital form. Abbreviated DAB. {di'j-əd-əl 'od-ē-ō 'brəd-kast-iŋ}

digital audio tape [COMPUT SCI] A magnetic tape on which sound is recorded and played back in digital form. Abbreviated DAT. {di'j-əd-əl 'od-ē-ō 'tāp}

digital camera [ELECTR] A television camera that breaks up a picture into a fixed number of pixels and converts the light intensity (or the intensities of each of the primary colors) in each pixel to one of a finite set of numbers. {di'j-əd-əl 'kam-rə}

digital channel [COMMUN] A transmission path that carries only digital signals. {di'j-əd-əl 'chan-əl}

digital chart [NAV] A navigational chart encoded in a computer-usable format and used, in combination with electronic devices, to produce a computer-generated video display which provides the navigator with an accurate pictorial presentation of the information normally gathered from a paper chart. Also known as electronic chart. {di'j-əd-əl 'chārt}

digital circuit [ELECTR] A circuit designed to respond at input voltages at one of a finite number of levels and, similarly, to produce output voltages at one of a finite number of levels. {di'j-əd-əl 'sər-kət}

digital circuit multiplication equipment [COMMUN] Equipment that uses digital compression techniques to increase the capacity of digital satellite and cable links carrying voice, facsimile, and voice-frequency modem traffic. {di'j-əd-əl 'sər-kət mäl-tə-plə'kə-shən i'kwip-mənt}

digital communications [COMMUN] System of telecommunications employing a nominally discontinuous signal that changes in frequency, amplitude, time, or polarity. {di'j-əd-əl kə'myü-nə'kə-shənz}

digital comparator [ELECTR] A comparator circuit operating on input signals at discrete levels. Also known as discrete comparator. {di'j-əd-əl kəm'par-əd-ər}

digital computer [COMPUT SCI] A computer operating on discrete data by performing arithmetic and logic processes on these data. {di'j-əd-əl kəm'pyüt-ər}

digital control [CONT SYS] The use of digital or discrete technology to maintain conditions in operating systems as close as possible to desired values despite changes in the operating environment. {di'j-əd-əl kən'tröl}

digital converter [ELECTR] A device that converts voltages to digital form; examples include analog-to-digital converters, pulse-code modulators, encoders, and quantizing encoders. {di'j-əd-əl kən'verd-ər}

digital counter [ELECTR] A discrete-state device (one with only a finite number of output conditions) that responds by advancing to its next output condition. {di'j-əd-əl 'kaunt-ər}

digital data [COMPUT SCI] Data that are electromagnetically stored in the form of discrete digits. {di'j-əd-əl 'dad-ə}

Digital Data Broadcast System [NAV] A system that will provide information aiding air-traffic control; digital data to aircraft over vortac channels will carry information on the geographic location, elevation, magnetic variation, and related data of the vortac station being received. Abbreviated DDBS. {di'j-əd-əl 'dad-ə 'brəd-kast 'sis-təm}

digital data modulation system [COMMUN] A digital communications system in which the information source consists of a finite number of discrete messages which are coded into a sequence of waveforms or symbols, each one selected from a specified and finite set. {di'j-əd-əl 'dad-ə mäj-ə'lə-shən 'sis-təm}

digital data recorder [COMPUT SCI] Electronic device that converts continuous electrical analog signals into number (digital) values and records these values onto a data log via a high-speed typewriter. {di'j-əd-əl 'dad-ə ri'kōrd-ər}

digital data service [COMMUN] A telephone communication system developed specifically for digital data, using existing local digital lines combined with data-under-voice

digital data service

605

DIGenea


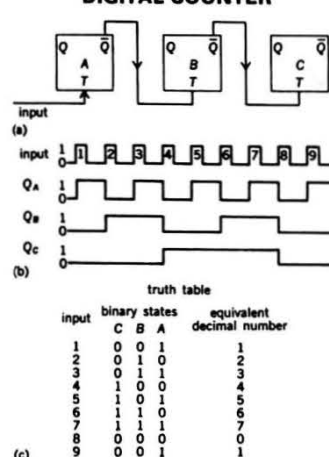


Diagram of an adult digenetic trematode. (From R. M. Cable, *An Illustrated Laboratory Manual of Parasitology*, Burgess, 1940)

DIGITAL COUNTER



An octal counter. T stands for trigger input, Q and Q represent output terminals, and A, B, and C identify different flip-flop stages. (a) Three successive flip-flop stages. (b) Input signal and Q-terminal states of each flip-flop. (c) Truth table.

Pacific time [ASTRON] The time for a given time zone that is based on the 120th meridian and is the eighth zone west of Greenwich. Also known as Pacific Standard Time. { pə'sif-ik 'tīm }

Pacific-type continental margin [GEOL] A continental margin typified by that of the western Pacific where oceanic lithosphere descends beneath an adjacent continent and produces an intervening island arc system. { pə'sif-ik 'tɪp kənt-ən'ent-əl 'mār-jən }

Pacinian corpuscle [NEUROSCI] An encapsulated lamellar sensory nerve ending that functions as a kinesthetic receptor. { pə'chin-ē-ən 'kɔr-pə-səl }

pack [COMPUT SCI] To reduce the amount of storage required to hold information by changing the method of encoding the data. [IND ENG] To provide protection for an article or group of articles against physical damage during shipment; packing is accomplished by placing articles in a shipping container, and blocking, bracing, and cushioning them when necessary, or by strapping the articles or containers on a pallet or skid. [MIN ENG] 1. A pillar built in the waste area or roadside within a mine to support the mine roof; constructed from loose stones and dirt. 2. Waste rock or timber used to support the roof or underground workings or used to fill excavations. Also known as fill. [OCEANOGR] See pack ice. [ORD] Part of a parachute assembly in which the canopy and shroud lines are folded and carried. Also known as pack assembly. { pak }

package [COMPUT SCI] A program that is written for a general and widely used application in such a way that its usefulness is not impaired by the problems of data or organization of a particular user. { pak-ij }

packaged circuit See rescap. { 'pak-ijd 'sər-kət }

packaged magnetron [ELECTR] Integral structure comprising a magnetron, its magnetic circuit, and its output matching device. { 'pak-ijd 'mag-nə-trən }

package freight [IND ENG] Freight shipped in lots insufficient to fill a complete car; billed by the unit instead of by the carload. { 'pak-ij frāt }

package power reactor [NUC PHYS] A small nuclear power plant designed to be crated in packages small enough for transportation to remote locations. { 'pak-ij 'paʊ-ər rē-ak-tər }

packaging [ELEC] The process of physically locating, connecting, and protecting devices or components. { 'pak-ə-jɪŋ }

packaging density [ELECTR] The number of components per unit volume in a working system or subsystem. { 'pak-ə-jɪŋ den-səd-ē }

pack artillery [ORD] Artillery weapons designed for transport in sections by animals or delivery by parachute; the weapon and carriage are partially disassembled for transport and reassembled for firing from ground positions. { 'pak ar'til-ə-rē }

pack assembly See pack. { 'pak ə-sem-blē }

pack builder [MIN ENG] 1. One who builds packs or pack walls. 2. In anthracite and bituminous coal mining, one who fills worked-out rooms, from which coal has been mined, with rock, slate, or other waste to prevent caving of walls and roofs, or who builds rough walls and columns of loose stone, heavy boards, timber, or coal along haulageways and passageways and in rooms where coal is being mined to prevent caving of roof or walls during mining operations. Also known as packer; pillar man; timber packer; waller. { 'pak bɪld-ər }

pack carburizing [MET] A method of surface hardening of steel in which parts are packed in a steel box with the carburizing compound and heated to elevated temperatures. { 'pak 'kär-bə-rɪz-ɪŋ }

packed bed [CHEM ENG] A fixed layer of small particles or objects arranged in a vessel to promote intimate contact between gases, vapors, liquids, solids, or various combinations thereof; used in catalysis, ion exchange, sand filtration, distillation, absorption, and mixing. { 'pakt 'bed }

packed decimal [COMPUT SCI] A means of representing two digits per character, to reduce space and increase transmission speed. { 'pakt 'des-məl }

packed file [COMPUT SCI] A file that has been encoded so that it takes up less space in storage. Also known as compressed file. { 'pakt 'fɪl }

packed tower [CHEM ENG] A fractionating or absorber tower filled with small objects (packing) to bring about intimate contact between rising fluid (vapor or liquid) and falling liquid. { 'pakt 'taʊ-ər }

packed tube [CHEM ENG] A pipe or tube filled with high-heat-capacity granular material; used to heat gases when tubes are externally heated. { 'pakt 'tʌb }

packer [ENG] A device that is inserted into a hole being grouted to prevent return of the grout around the injection pipe. [MIN ENG] See pack builder. [PETRO ENG] See production packer. { 'pak-ər }

packer fluid [PETRO ENG] Fluid inserted in the annulus between the tubing and casing above a packer in order to reduce pressure differentials between the formation and the inside of the casing and across the packer. { 'pak-ər flʊ-əd }

packer test [PETRO ENG] A pressure test of a sealed zone in a well. { 'pak-ər ,test }

packet [BIOL] A cluster of organisms in the form of a cube resulting from cell division in three planes. [COMMUN] A short section of data of fixed length that is transmitted as a unit. [PHYS] See wave packet. { 'pak-ət }

packet gland [INV ZOO] A cluster of gland cells opening through the epidermis of nemertines. { 'pak-ət ,gland }

packet switching See packet transmission. { 'pak-ət ,swɪtʃ-ɪŋ }

packet transmission [COMMUN] Transmission of standardized packets of data over transmission lines rapidly by networks of high-speed switching computers that have the message packets stored in fast-access core memory. Also known as packet switching. { 'pak-ət tranz-mɪʃ-ən }

pack hardening [MET] A process of heat treating in which the workpiece is packed in a metal box together with carbonaceous material; carbon penetration is proportional to the length of heating; after treatment the workpiece is reheated and quenched. { 'pak ,hɑrd-ən-ɪŋ }

pack ice [OCEANOGR] Any area of sea ice, except fast ice, composed of a heterogeneous mixture of ice of varying ages and sizes, and formed by the packing together of pieces of floating ice. Also known as ice canopy; ice pack; pack. { 'pak ,ɪs }

packing [CRYSTAL] Arrangement of atoms or ions in a crystal lattice. [ENG] See stuffing. [ENG ACOUS] Excessive crowding of carbon particles in a carbon microphone, produced by excessive pressure or by fusion particles due to excessive current, and causing lowered resistance and sensitivity. [GEOL] The arrangement of solid particles in a sediment or in sedimentary rock. [GRAPHICS] Paper used as a layer under the image or impression cylinder in letterpress printing or under the plate or blanket in lithographic printing in order to produce suitable pressure. [MET] In powder metallurgy, a material in which compacts are embedded during presintering or sintering operations. { 'pak-ɪŋ }

packing density [COMPUT SCI] The amount of information per unit of storage medium, as characters per inch on tape, bits per inch or drum, or bits per square inch in photographic storage. [ELECTR] The number of devices or gates per unit area of an integrated circuit. [GEOL] A measure of the extent to which the grains of a sedimentary rock occupy the gross volume of the rock in contrast to the spaces between the grains; equal to the cumulative grain-intercept length along a traverse in a thin section. { 'pak-ɪŋ den-səd-ē }

packing fraction [NUC PHYS] The quantity $(M - A)/A$, where M is the mass of an atom in atomic mass units and A is its atomic number. { 'pak-ɪŋ ,frak-shən }

packing house [FOOD ENG] 1. A food processing plant generally requiring the use of refrigeration. 2. A building in which livestock are slaughtered and processed, and the meat products and by-products are packed. { 'pak-ɪŋ ,haʊs }

packing house pitch [MATER] Dark-brown to black by-product residue from manufacturing soap and candle stock or from refining vegetable oils, refuse, or wool grease; soluble in naphtha and carbon disulfide; used to make paints, varnishes, and tar paper, and in marine caulking and waterproofing. Also known as fatty-acid pitch. { 'pak-ɪŋ ,haʊs ,pɪtʃ }

packing index [CRYSTAL] The volume of ion divided by the volume of the unit cell in a crystal. { 'pak-ɪŋ ,ɪn-deks }

packing proximity [GEOL] In a sedimentary rock, an estimate of the number of grains that are in contact with adjacent grains; equal to the total percentage of grain-to-grain contacts along a traverse measured on a thin section. { 'pak-ɪŋ ,prɒk-sɪm-əd-ē }

packing radius [CRYSTAL] One-half the smallest approach distance of atoms or ions. { 'pak-ɪŋ ,ræd-ē-əs }

HARGRAVE'S COMMUNICATIONS DICTIONARY

Frank Hargrave



**IEEE
PRESS**

The Institute of Electrical and Electronics Engineers, Inc., New York

This book and other books may be purchased at a discount from the publisher when ordered in bulk quantities. Contact:

IEEE Press Marketing
Attn: Special Sales
445 Hoes Lane
P.O. Box 1331
Piscataway, NJ 08855-1331
Fax: +1 732 981 9334

For more information about IEEE Press products, visit the IEEE Online Store & Catalog: <http://www.ieee.org/store>.

© 2001 by the Institute of Electrical and Electronics Engineers, Inc.
3 Park Avenue, 17th Floor, New York, NY 10016-5997.

All rights reserved. No part of this book may be reproduced in any form, nor may it be stored in a retrieval system or transmitted in any form, without written permission from the publisher.

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

ISBN 0-7803-6020-6
IEEE Order No. PC5869

Library of Congress Cataloging-in-Publication Data

Hargrave, Frank.

Hargrave's communications dictionary.

/ Frank Hargrave.

p. cm.

Includes index.

ISBN 0-7803-6020-6

1. Telecommunication—Dictionaries. I. Title.

TK5102.H37 2000

621.382'03—dc21

00-061416


DAMA

DAMA (1) An acronym from Demand Assigned Multiple Access. (2) An acronym from Data Assigned Multiple Access.

damped response In linear systems including circuits, networks, and mechanical devices, a term describing the system output (response) to an abrupt input stimulus (a step or impulse for example). There are three classifications of response—*critically damped*, *underdamped*, and *overdamped*.

- **Critically damped**—In response to an abrupt input stimulus, the output of a critically damped system will change in the minimum time possible *without* overshoot or ringing. It is the boundary between an underdamped and overdamped response.
- **Underdamped**—In response to an abrupt input stimulus, the output of an underdamped system has an overshoot and possible ringing. Overshoot is the condition of a signal to exceed the final value when changing from one state to another.
- **Overdamped**—As with a critically damped system, an overdamped system will not have either overshoot or ringing. However, the system output response to an abrupt input signal change will be slower than that of the critically damped system.

A comparison of the three damping conditions is shown in the accompanying figure.



damping The progressive diminishing of specified values characteristic of a phenomenon with respect to time, e.g., the progressive decay in the amplitude of the free oscillations of a circuit with time. See also *damped response*.

damping factor (ζ) A number expressing the ratio of the actual damping to the damping of a critically damped second-order linear system (or subsystem). *Damping factors* less than one are underdamped, while those greater than one are overdamped. See also *damped response*.

DAN An acronym from Departmental Area Network.

DAP (1) An acronym from Data Access Protocol. (2) An acronym from Directory Access Protocol. (3) An acronym from Document Application Profile.

dark current In electro-optics, the quiescent current that flows in a photosensitive device when there is no incident radiation.

dark fiber An unused optic fiber. A fiber optic cable not carrying a signal; that is, there is no light energy in the fiber.

DARPA An acronym from Defense Advanced Research Projects Agency.

DAS (1) An acronym from Dual Attachment Station. A device on an FDDI fiber optic ring to which the two rings are attached. (2) An acronym from Dynamically Assigned Socket. (3) An acronym from Disk Array Station.

135

data bits

DASD An acronym from Direct Access Storage Device (pronounced "DAZ-dee").

DASS (1) An acronym from Distributed Authentication Security Service. (2) An acronym from Direct Access Secondary Storage. Storage facilities that are online but slower than the mainstay hard disk drive.

DASS n A British acronym from Digital Access Signaling System. DASS 1 was the original British Telecom ISDN signaling scheme developed for both single line and multiline Integrated Digital Access (IDA) applications (although it was used only with single lines). DASS 2, also developed by British Telecom, is a message-based signaling system that follows the ISO model and provides multiline access to the British Telecom network.

DassIII A message-based signaling system, following the ISO model, developed by British Telecom (BT), to provide multiline integrated digital access (IDA) interconnection to the BT network.

DAT (1) An acronym from Digital Audio Tape. (2) An acronym from Duplicate Address Test.

data A representation of a collection of facts, concepts, instructions, or information to which meaning has been assigned. The representation may be analog, digital, or any symbolic form suitable for storage, communication, interpretation, or processing by human or automatic means.

"Data" is the plural of the Latin *datum*, meaning one item of information. To be correct, a single item should be called a datum and more than one should be called *data*, i.e., "one datum is . . ." and "two data are . . ."

data above voice (DAV) See *data over voice*.

data access arrangement (DAA) See *DAA*.

Data Access Language (DAL) In Macintosh-based client-server environments, an extension to the Structured Query Language (SQL) database language intended to provide uniform access to any database that supports SQL.

data access manager (DAM) In the Apple Computer's System 7 operating system software for Macintoshes, *DAM* is a built-in capability for accessing databases on a network. The *DAM* mediates between an application and the database being accessed. It uses database extensions to communicate with the database. These are database-specific system files that contain the commands necessary to interact with a particular database.

data attribute A characteristic of a data element such as length (number of bits or bytes), method of representation (fixed point, floating point, alphanumeric), or value.

data bank A set of data related to a given subject and organized in such a way that it can be accessed and retrieved by local and/or remote users. The databank's characteristics and attributes may include:

- Information on a single subject or multiple subjects.
- Any rational organization (random, sequential, and so on).
- More than one database.
- More than one data bank in order to achieve a complete database.
- May be geographically distributed.

data base See *database*.

data bits In asynchronous communications, the group of 5, 6, 7, or 8 bits following the start bit (7 or 8 bits are most commonly used). These bits represent a single character or symbol for transmission. Following the *data bits* is an optional parity bit and 1, 1.5, or 2 stop bits.

SECOND EDITION

COMPREHENSIVE

DICTIONARY

OF
ELECTRICAL
ENGINEERING

EDITOR-IN-CHIEF
Phillip A. Laplante



Taylor & Francis
Taylor & Francis Group

Boca Raton London New York Singapore

CRC PRESS, a Taylor & Francis title, part of the Taylor and Francis Group.

Published in 2005 by
CRC Press
Taylor & Francis Group
6000 Broken Sound Parkway NW, Suite 300
Boca Raton, FL 33487-2742

© 2005 by Taylor & Francis Group, LLC
CRC Press is an imprint of Taylor & Francis Group

No claim to original U.S. Government works
Printed in the United States of America on acid-free paper
10 9 8 7 6 5 4 3 2 1

International Standard Book Number-10: 0-8493-3086-6 (Hardcover)
International Standard Book Number-13: 978-0-8493-3086-5 (Hardcover)
Library of Congress Card Number 2004058572

This book contains information obtained from authentic and highly regarded sources. Reprinted material is quoted with permission, and sources are indicated. A wide variety of references are listed. Reasonable efforts have been made to publish reliable data and information, but the author and the publisher cannot assume responsibility for the validity of all materials or for the consequences of their use.

No part of this book may be reprinted, reproduced, transmitted, or utilized in any form by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying, microfilming, and recording, or in any information storage or retrieval system, without written permission from the publishers.

For permission to photocopy or use material electronically from this work, please access www.copyright.com (<http://www.copyright.com/>) or contact the Copyright Clearance Center, Inc. (CCC) 222 Rosewood Drive, Danvers, MA 01923, 978-750-8400. CCC is a not-for-profit organization that provides licenses and registration for a variety of users. For organizations that have been granted a photocopy license by the CCC, a separate system of payment has been arranged.

Trademark Notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

Library of Congress Cataloging-in-Publication Data

Comprehensive dictionary of electrical engineering / editor-in-chief Phillip A. Laplante.-- 2nd ed.
p. cm.

ISBN 0-8493-3086-6 (alk. paper)

1. Electric engineering--Dictionaries. I. Title: Electrical engineering. II. Laplante, Phillip A.

TK9.C575 2005

621.3'03--dc22

2004058572



Taylor & Francis Group
is the Academic Division of T&F Informa plc.

Visit the Taylor & Francis Web site at
<http://www.taylorandfrancis.com>

and the CRC Press Web site at
<http://www.crcpress.com>

amplitude-modulated link

signal and reduce the frequency depression (especially in Meachem-bridge oscillator with crystal) of the main harmonic by higher harmonics (van der Pol effect). Three types of circuits are used:

1. An element of large inertia (tungsten lamp, thermistor) is included in the circuit at a point where it can change the magnitude of feedback, but not affect the frequency.
2. A controlled resistor (usually an FET operating in a triode regime) that is also part of the feedback circuit (the DC control signal is obtained with a rectifier and a filter of large time constant).
3. An automatic gain control circuit where the DC control signal obtained from a rectifier and filter is used to change the bias of oscillator active element.

amplitude-modulated link a transmitter–receiver system that utilizes amplitude-modulation for the transmission of signal frequencies.

amplitude-shift keying (ASK) a modulation technique in which each group of source bits determines the amplitude of the modulated carrier.

AMPS See advanced mobile phone system.

AMR See automated meter reading.

analog See analog signal, analog data.

analog data data represented in a continuous form with respect to continuous time, as contrasted with digital data represented in a discrete (discontinuous) form in a sequence of time instant.

analog multiplier a device or a circuit that generates an analog output signal that is proportional to the product or multiplication of two analog input signals.

analog optical computing optical computing that involves two-dimensional analog operations such as correlation and complex spatial frequency filtering primarily based on the property of the lens

to perform two-dimensional Fourier transform. In analog optical computing, operations to be performed are matched with and based on already known optical phenomena.

analog signal a signal represented in a continuous form with respect to continuous time, as contrasted with digital signal represented in a discrete (discontinuous) form in a sequence of time instant. *See also* analog data.

analog signal conditioning an interface between the sensor or transducer output, which represents an analog or physical world, and the analog-to-digital converter.

analog-to-digital A/D conversion a method by which a continuously varying signal (voltage) is sampled at regularly occurring intervals. Each sample is quantized to a discrete value by comparisons to preestablished reference levels. These quantized samples are then formatted to the required digital output (e.g., binary pulse code words). The A/D converter is “clocked” to provide updated outputs at regular intervals. In order not to lose any baseband information, sampling must occur at a rate higher than twice the highest incoming signal frequency component. *See also* Nyquist rate.

analog-to-digital A/D converter a device that changes an analog signal to a digital signal of corresponding magnitude. This device is also called an encoder, ADC, or A/C converter.

analysis filter a filter in the analysis section of a sub-band analysis and synthesis system.

analysis-by-synthesis coding refers to the class of source coding algorithms where the coding is based on parametric synthetization of the source signal at the encoder. The synthesized signal is analyzed, and the parameters that give the “best” result are chosen and then transmitted (in

NEWTON's TELECOM DICTIONARY

STAY INFORMED

To be alerted by email to updates and corrections
go to www.cmpbooks.com/newton

CMPBooks
San Francisco

NEWTON'S TELECOM DICTIONARY

copyright © 2003 Harry Newton

email: Harry@HarryNewton.com

personal web site: www.HarryNewton.com

business web site: www.TechnologyInvestor.com

All rights reserved under International and Pan-American Copyright conventions,
including the right to reproduce this book or portions thereof in any form whatsoever.

Published by CMP Books

An imprint of CMP Media LLC

Main office: CMP Books, 600 Harrison St., San Francisco, CA 94107 USA
Phone: 415-947-6615; Fax: 415-947-6015

Sales office: CMP Books, 12 West 21 Street, New York, NY 10010
Phone 917-305-3333; Fax 212-206-0387

www.cmpbooks.com

Email: books@cmp.com



CMP

United Business Media

For individual orders, and for information on special discounts for quantity orders,
please contact:

CMP Books Distribution Center, 6600 Silacci Way, Gilroy, CA 95020

Tel: 1-800-500-6875 or 408-848-3854; Fax: 408-848-5784

Email: cmp@rushorder.com; Web: www.cmpbooks.com

Distributed to the book trade in the U.S. by:

Publishers Group West, 1700 Fourth Street, Berkeley, California 94710

Distributed in Canada by:

Jaguar Book Group, 100 Armstrong Avenue, Georgetown, Ontario M6K 3E7 Canada

Printed in the United States of America

This book is also sold through www.Amazon.com, www.Fatbrain.com and
www.BarnesAndNoble.com

Distributed to the book trade in the U.S. and Canada by Publishers Group West
1700 Fourth St., Berkeley, CA 94710

Fax: 408-848-5784

cmp@rushorder.com

ISBN Number 1-57820-307-4

March 2003

Nineteenth Edition

Matt Kelsey, Publisher

Ray Horak, Senior Contributing Editor

Saul Roldan, Cover Artist

Lisa Giaquinto, Project Manager

Brad Greene, Text Layout

changes in the environment in which they must operate. See also Analog Computer.

Analog Recording System of recording in which music is converted into electrical impulses that form "patterns" in the grooves of phonograph record masters or in the oxide particles of master tapes representing (or analogous to) musical waveforms.

Analog Semiconductor Analog semiconductors are essentially the "translators" between the wave-form world of man (light, heat, pressure, and sound all move in waves) and the digital world ("ones" and "zeros") of computers. Analog semiconductors act as amplifiers in strengthening a weak signal, as converters to turn a signal from wave-form to digital and back again, and as voltage regulators, stepping down a signal from higher to lower power, also called power management.

Analog Signal A signal in the form of a continuous wave varying in step with the actual transmitted information; attempts to transmit an exact replica of the inputted signal down a communications channel. See Analog and all the various definitions starting with Analog.

Analog Switch Telephone switching equipment that switches signals without changing the analog form of the original phone call. The major form of analog switching is circuit switching. Most switching is now done digitally.

Analog Synchronization A synchronization control system in which the relationship between the actual phase error between clocks and the error signal device is a continuous function over a given range.

Analog Transmission A way of sending signals — voice, video, data — in which the transmitted signal is analogous to the original signal. In other words, if you spoke into a microphone and saw your voice on an oscilloscope and you took the same voice as it was transmitted on the phone line and threw that signal onto the oscilloscope, the two signals would look essentially the same. The only difference would be that the electrically transmitted signal would be at a higher frequency. Most transmission is now done digitally.

Analog Video Signals represented by an infinite number of smooth transitions between video levels. TV signals are analog. By contrast, a digital video signal assigns a finite set of levels. Because computer signals are digital, analog video must be converted into a digital form before it can be shown on a computer screen.

Analog Wireless The dominant radio transmission standard in the United States; also called AMPS.

Analogous An English/European way of spelling analog, which is the correct North American spelling. See Analog.

Anamorphic Unequally scaled in vertical and horizontal dimensions.

ANC All Number Calling. The dialing plan used in telephone networks. Consisting of all numbers, ANC replaced the old U.S. system which consisted of two letters and five numbers (2L + 5N). In other words, the GR (Greenwood) exchange became 47, PA (Pennsylvania) became 72, and UL (Ulysses) became 85. Remember the Glenn Miller hit, "Pennsylvania 6-5000?" Those old exchanges were charming, often reflecting the character of the communities to which they were assigned. However, and as the number of telephone numbers grew, we ran out of alpha prefixes that included the first two letters of meaningful words. Eventually, we had to use the 1 on the dial in the second position of the prefix; there are no letters associated with that number. Hence, All Number Calling.

ANCARA Advanced Networked Cities And Regions Association. A formal association of cities and regions exploring advanced uses of information technology. ANCARA was founded in 1996 by the regions/cities of Eindhoven (The Netherlands), Kansai (Japan), Orlando (Florida), Silicon Valley (California), Singapore, and Stockholm (Sweden). The intent is to accelerate the development of the Global Information Infrastructure (GII). www.ancara.nl. See also GII.

Ancestor Node An ATM term. A logical group node that has a direct parent relationship to a given node (i.e., it is the parent of that node, or the parent's parent.)

Anchor A hyperlinked word or group of words. An anchor is the same as a hyperlink — the underlined words or phrases you click on in World Wide Web documents to jump to another screen or page. The word anchor is used less often than hyperlink, but it maintains the seafaring theme of navigating and surfing the Net. See also Hyperlink.

Anchorage Accord A milestone ATM Forum document (April 12, 1998) so named because of the meeting location, the Anchorage Accord outlines which versions of ATM Forum specifications vendors should implement. ATM Forum specifications comprise approximately 60 baseline specifications for successful market entry of ATM products and services. Included are Broadband InterCarrier Interface (BICI), Interim Local Management Interface (ILMI), LAN Emulation (LANE), network management, Private Network Node Interface (PNNI), signaling, SMDS (Switched Multimegabit Data Service) and IP (Internet

Protocol) over ATM, traffic management, and a number of physical interfaces. The accord also limits the conditions under which specifications are revised in order to cut down on future confusion. See also ATM Forum.

Anchor System A cellular term. An anchor system is the system that maintains the connection to the PSTN (Public Switched Telephone Network) during the process of call handoff from one cell site to another. In this context, a "system" comprises all of the cellular carrier's MSCs (Mobile Switching Centers) serving a particular geographic area, and all of the cell sites supported by the MSCs. See also MSC.

Ancillary Charges Charges for supplementary services comprised of optional features, which may consist of both non-recurring and monthly charges.

AND Automatic Digital Network Dialing. A digital private line service that transmits voice, data, video and other digital signals.

AND Gate A digital device which outputs a high state if either of its inputs are high.

AND Logic Gate A type of logic gate which uses AND logic. The output of an AND logic gate would consider the first and second input.

And Statement See Not Statement.

Anechoic Chamber A perfectly quiet room. A room in which sound or radio waves do not reflect off the walls. An anechoic chamber is the only place in which a speakerphone will work perfectly. The more a room resembles an anechoic chamber — i.e. lots of drapes, plush carpet, etc. — the better a speakerphone will work.

Anemometer A device which measures wind speed and direction.

Angel Investor in an early stage technology start-up. Typically an angel invests when the company is little more than an idea, a simple business plan and several management people, but rarely a full management team.

Angle Bracket The term for these two brackets < and >. These two brackets have major use in the HTML language. See HTML.

Angle Modulation Modulation in which phase angle or frequency of a sine wave carrier is varied.

Angle of Arrival AOA. A class of Position Determination Technology in which a mobile radio unit's position is calculated based on the direction of its transmitted signal measured from two or more receiving sites; also known as triangulation. Employed in certain wireless E9-1-1 solutions. See also Time Difference of Arrival.

Angle Of Deviation In fiber optics, the net angular deflection experienced by a light ray after one or more refractions or reflections. The term is generally used in reference to prisms, assuming air interfaces. The angle of deviation is then the angle between the incident ray and the emergent ray.

Angle of Incidence The angle between an incident ray and the normal to a reflecting or refracting surface.

Angled End An optical fiber whose end is deliberately polished at an angle to reduce reflections.

Angry fruit salad A terrible visual interface on a web site or a software screen that has far too many colors.

Angstrom One ten-millionth of a millimeter. Angstroms are primarily used to express electromagnetic wavelengths, including and particularly optical wavelengths. It is also called an Angstrom unit. It is named after Anders Jonas Angstrom, a Swedish astronomer and physicist, who lived 1814-74. It is said that the 128-bit addressing scheme of IPv6 provides enough unique IP addresses to theoretically provide 1,500 per square angstrom of the earth's surface.

Angular circumference The measurement of the amount of bend in a fiber-optic cable.

Angular Misalignment Loss The optical power loss caused by angular deviation from the optimum alignment of source to optical fiber — fiber-to-fiber, or fiber-to-detector.

ANI Automatic Number Identification. ANI provides for the transmission through the network of the BN (Billing Number), versus the telephone number, of the originating party (i.e., the calling person, also called party in the phone business). ANI originally was intended exclusively for the use of the long distance and local phone carriers for billing purposes. ANI information is sent through the network, from the originating central office, through all intermediate tandem offices, to the terminating central office. The information originally was sent over analog trunks in the form of DTMF (Dual Tone MultiFrequency) signals, although contemporary networks usually pass the information through the digital SS7 (Signaling System 7) network. For some years, ANI has been available to end user organizations, as well. In order to gain access to ANI data, you must have a "trunk side" con-

are used to impress people. "This page is kinda dull. Maybe a little dancing baloney will help?" This definition courtesy Wired Magazine.

Dancing frog A problem or image on your computer screen that disappears just as soon as you try to show it to someone else. The same thing seems to happen with automobiles when you take a normally troublesome car in for a checkup with the mechanic.

DAP Directory Access Protocol. The protocol used between a Directory User Agent (DUA) and Directory System Agent (DSA) in an X.500 directory system. See X.500 and LDAP.

Dark Optical fiber through which no light is currently being transmitted. See Dark Fiber.

Dark Current The flow of electricity through the diode in a photodiode when no light is present. Photodiodes are often used as light-sensitive switches. When light hits them, they turn on. Here's a more technical explanation: Dark current is the induced current that exists in a reversed biased photodiode in the absence of incident optical power. It is better understood to be caused by the shunt resistance of the photodiode. A bias voltage across the diode (and the shunt resistance) causes current to flow in the absence of light. See also Dark Fiber.

Dark Fiber Optical fiber through which no light is transmitted and which, therefore, no signal is being carried. Generally speaking, a dark fiber is one of many fibers contained within a cable. Carriers commonly deploy a large number of fibers (432 is a common number) at any given time, since the incremental cost is quite modest compared to pulling them one at a time as the need arises. In fact, a carrier often has little choice, as the right of way may be granted once, and only once. The fibers that the carrier is using immediately are "lit," and those that currently are unused are left "dark." The dark fiber is available for future use. Sometimes dark fiber is sold by a carrier without the accompanying transmission electronics. The customer, which may be either an end user organization or another carrier, is expected to light it up with his own electronics. See also Dark Current, Dim Fiber and Lit Fiber.

Dark Side At Apple trade shows, people who use Windows machines are known as being on the Dark Side.

Dark Swap Round-trip commodity trading of unused broadband — so-called dark fiber — among providers. The technique creates the appearance of trade activity. An unscrupulous carrier can book the swapping as revenue and thus make his financials look better to investors in the stockmarket. It happened in the late 1990s and early 2000s. And by the time you read this, some of the creative executives who thought this up should be sitting firmly in jail.

Dark Wavelength A Dense Wavelength Division Multiplexing (DWDM) term. Dark wavelength refers to a virtual channel in a fiber optic system utilizing DWDM. Each virtual channel is supported through a specific wavelength of light, with many such channels riding over the same fiber. Once the fiber system is deployed and the DWDM equipment is activated, some of the wavelengths may be activated immediately and others may be left dark for future needs. Such a fiber system is called "Dim Fiber," as it's neither completely dark, nor fully lit. When the need arises, those dark wavelengths are lit up. See also DWDM, Fiber, Optical Fiber and SONET.

DARPA Defense Advanced Research Projects Agency. Formerly called ARPA, it is a US government agency that funded research and experimentation with the ARPANET and later the Internet. The group within DARPA responsible for the ARPANET is ISTO (Information Systems Techniques Office), formerly IPTO (Information Processing Techniques Office). DARPA had sponsored research in the 1960s and the 1970s to create a computer network that could survive a nuclear detonation. See also DARPA Internet, IAB, IETF and Internet.

DARPA Internet World's largest internetwork, linking together thousands of networks around the world. Sponsored by U.S. Defense Advanced Research Projects Agency. Now called DARPA Internet. See next definition.

DARPA Internet Defense ARPANET. Also known as DARPA Internet. In 1983 the ARPANET was officially split into DARPA Internet and MILNET. World's largest internetwork, linking together thousands of networks around world. Sponsored by U.S. Defense Advanced Research Projects Agency. DARPA Internet was the beginnings of the Internet. See Internet.

DARS Digital Audio Radio System. Also known as DAB (Digital Audio Broadcasting) outside the U.S. Proposed satellite-delivered audio/radio systems, similar to DBS (Direct Broadcast System) TV systems, which have been enormously successful in competition with CATV. DARS has been debated by the FCC and the ITU-R since the initial application by CD Radio Inc. in 1990. Assuming that the FCC and ITU-R eventually agree on frequency assignments (and they now have), you may want to make room for one more satellite dish on your rooftop or on your car.

DAS Tape A cellular term. The magnetic tape that is used at the MTSO to record traf-

fic statistics and call billing information. This tape is sent to a third-party 'billing-house' where the actual billing of the subscribers is done.

DASD Direct Access Storage Device. Any on-line data storage device. Usually refers to a magnetic disk drive, because optical drives and tape are considered too slow to be direct access devices. Pronounced DAZ-dee. The term is said to have been invented by IBM.

DASS Direct Access Secondary Storage. Same as near-line: storage on pretty-fast storage devices (e.g., rewritable optical) that are less expensive than hard drives but faster than off-line devices.

DASS1 Digital Access Signaling. A British term. The original British Telecom (BT) ISDN signalling developed for both single line and multi-line Integrated Digital Access but used in the BT ISDN pilot service for single line IDA only.

DASS2 Digital Access Signaling System No. 2. A British Term. A message-based signalling system following the ISO-based model developed by British Telecom to provide multi-line IDA interconnection to the BT network.

DAT Digital Audio Tape used to identify a type of digital tape recorder and player as well as the tape cassette. DAT tape machines record music that is much crisper, and free of the hisses and pops that mar traditional analog recordings. The drawback with DAT tape machines is they require considerable tape to store music digitally. In a DAT machine, the music is recorded by sampling the music 48,000 times each second. Each of those samples is represented by a number that is written as a 16-digit string of zeros and ones. There are two such signals, once for each stereo channel, meaning that storing a single second of music requires about 1.5 million bits. On top of that, extra bits are added to allow the system to mathematically correct errors and help the machine automatically find a particular song on the tape. All together, according to Andrew Pollack writing in the New York Times, a single second of music on a digital audio tape requires 2.8 million bits. But compression techniques are cutting down the amount of information required to be recorded.

Data This is AT&T Bell Labs' definition: "A representation of facts, concepts or instructions in a formalized manner, suitable for communication, interpretation or processing." Typically anything other than voice.

Data Abstraction A term in object-oriented programming. An object is sometimes referred to as an instance of an abstract data type or class. Abstract data types are constructed using the built-in data types supported by the underlying programming language, such as integer and date. The common characteristics (both attributes and methods) of a group of similar objects are collected to create a new data type or class. Not only is this a natural way to think about the problem domain, it is a very efficient way to write programs. Instead of individually describing several dozen instances, the programmer describes the class once. Once identified, each instance is complete with the exception of its instance variables. The instance variables are associated with each instance, i.e., each object; methods exist only with the classes. See Object Oriented Programming.

Data Access Arrangement DAA. Equipment that allows you to attach your data equipment to the nation's phone system. At one stage, DAAs were required by FCC "law." Now, their limited functions are built into directly attached devices, such as terminals, computers, etc.

Data Access Point DAP. MCI computer that holds the number translation and call-routing information for 800 and Vnet services. These computers respond to inquiries from MCI switches on how to handle these calls.

Data Arrangement In public switched telephone networks, a single item or group of items present at the customer's premises, including all equipment that may affect the characteristics of the interface. An obsolete term. Historically, it came from the time when the phone industry insisted on an interface between its lines and equipment provided by others.

Data Attribute A characteristic of a data element such as length, value, or method of representation.

Data Bank A collection of data in one place. The data is not necessarily logically related, nor is it necessarily consistently maintained. See Database.

Data Base See Database, which is our preferred spelling.

Data Broadcasting A method of high speed data distribution for text and graphics which uses the spare capacity in the broadcasting television, cable and satellite transmission systems.

Data Bubble A new organization within BellSouth to provide high-speed digital services. No one seems to know why it's called "Data Bubble," except that someone inside BellSouth clearly thinks the term is cute.

Data Burst Burst transmission.

Digital Multiplex Hierarchy

ly were of this type before VGA models appeared. Digital monitors do not have as wide a range of color choices as analog types; digital EGA monitors, for example, can display just 16 colors out of a palette of 64.

Digital Multiplex Hierarchy An ordered scheme for the combining of digital signals by the repeated application of digital multiplexing. Digital multiplexing schemes may be implemented in many different configurations depending upon the number of channels desired, the signaling system to be used, and the bit rate allowed by the communication medium. Some currently available multiplexers have been designated as D1-, DS-, or M-series, all of which operate at T-carrier rates. Extreme care must be exercised when selecting equipment for a specific system to ensure interoperability, because there are incompatibilities among manufacturers' designs (and various nations' standards).

Digital Multiplexed Interface A ISDN PRI-like connection between a PBX and a computer, developed by AT&T.

Digital Multiplexer A device for combining digital signals. Usually implemented by interleaving bits, in rotation, from several digital bit streams either with or without the addition of extra framing, control, or error detection bits. In short, equipment that combines by time division multiplexing several signals into a single composite digital signal.

Digital Nervous System Coined by Bill Gates in 1997, the best definition of this term came from an interview between Gary Reiner, GE's chief information officer and a reporter from the Economist. According to the magazine, "Mr Reiner heads the company's most important initiative: 'digitising' as much of its business as possible. That not only means buying and selling most things online but, more importantly, setting up a digital nervous system that connects in real time anything and everything involved in the company's business: IT (Information Technology) systems, factories and employees, as well as suppliers, customers and products."

Digital Network A network in which the information is encoded as a series of ones and zeros rather than as a continuously varying wave — as in traditional analog networks. Digital networks have several major pluses over analog ones. First, they're "cleaner." They have far less noise, static, etc. Second, they're easier to monitor because you can measure them more easily. Third, you can typically pump more digital information down a communications line than you can analog information.

Digital Network Architecture. DNA. The data network architecture of Digital Equipment Corporation (DEC), now part of Compaq Corporation.

Digital Phase-Locked Loop A phase-locked loop in which the reference signal, the controlled signal, or the controlling signal, or any combination of these, is in digital form.

Digital Phase Modulation The process whereby the instantaneous phase of the modulated wave is shifted between a set of predetermined discrete values in accordance with the significant conditions of the modulating digital signal.

Digital Plastic A fancy term for buying goods and services on-line over the Internet using your credit card, possibly in conjunction with some verification of who you are from an independent certification authority.

Digital Port Adapter DPA. A device which provides conversion from the RS-449/422 interface to the more common interfaces of RS-232-C, V.35, WE-306 and others.

Digital Private Network Signaling System See DPNSS.

Digital Pulse Origination DPO. Equipment that sends dialed digits consisting of tones or pulses. It may be used at the central office end of a DID service connection.

Digital Pulse Termination DPT. Equipment that receives and processes dialed digits consisting of tones or pulses. It may be used at the customer end of a DID service connection.

Digital Radio Broadcasting DRB. Radio transmission intended for general reception in the form of discrete, integral values.

Digital Radio Concentrator System DRCS. A digital radio system which transmits data via a device which connects a number of circuits, which are not all used at once, to a smaller group of circuits for economy.

Digital Recording A system of recording by conversion of musical information into a series of pulses that are translated into a binary code intelligible to computer circuits and stored on magnetic tape or magnetic discs. Also called PCM - Pulse Code Modulation.

Digital Reference Signal DRS. A digital reference signal is a sequence of bits that represents a 1004-Hz to 1020-Hz signal.

Digital Selective Calling DSC. A synchronous system developed by the International Radio Consultative Committee (CCIR), used to establish contact with a station

or group of stations automatically by radio. The operational and technical characteristics of this system are contained in CCIR Recommendation 493.

Digital Sequence Spread Spectrum A wireless term. An RF (radio frequency) modulation technique, which uses algorithms to code transmissions in sequential channels and then decode them at the other end.

Digital Service Cross-Connect DSX. A termination/patch panel that lets DS1 and DS3 circuits be monitored by test equipment.

Digital Set-Top Box A device that hooks up to a TV and can collect, store, and display digitally compressed TV signals. See also Digital Cable Set Top Box.

Digital Signal A discontinuous signal. One whose state consists of discrete elements, representing very specific information. When viewed on an oscilloscope, a digital signal is "squared." This compares with an analog signal which typically looks more like a sine wave, i.e. curvy. Usually amplitude is represented at discrete time intervals with a digital value.

Digital Signal Cross-Connect DSX. Also known variously as a DACS (Digital Access Cross-Connect System) and a DCC (Digital Cross-Connect), a DSX is a device that is used to connect digital circuits together. A DSX-1 interconnects DS-1 (T-1 or E-1) circuits, as DSX-2 interconnects DS-2 (T-2 or E-2) circuits, and a DSX-3 interconnects DS-3 circuits (T-3 or E-3). Digital Signal Level DS-n. A hierarchical arrangement of digital signals used in North America beginning with DS-0 (64 Kbps) up to DS-4 (274 Mbps).

Digital Signal Processor A digital signal processor is a specialized semiconductor device or specialized core in a semiconductor device that processes very efficiently and in real time a stream of digital data that is sampled from analog signals ranging from voice, audio and video and from cellular and wireless to radio and television. As opposed to a general-purpose processor, a DSP is often designed to solve specific processing problems. A DSP architecture focuses on algorithmic efficiency and may use an instruction set that is more or less tailored toward the problem the DSP is solving. General purpose processors, on the other hand, may sacrifice algorithmic efficiency for general-purpose capability and push clock-speed to achieve performance. A DSP typically has much greater mathematical computational abilities than a standard microprocessor. In some applications, like wireless, PDAs and cell phones, constraints on power consumption require performance improvements other than faster clock speed. In other applications, like cellular base stations and high definition TV, where the number of channels or the high data rate require signal processing capabilities an order of magnitude greater than general purpose processors, a DSP that uses processing parallelism can provide much higher performance much more efficiently than even the fastest general-purpose processor. A DSP often performs calculations on digitized signals that were originally analog (e.g. voice or video) and then sends the results on. There are two main advantages of DSPs — first, they have powerful mathematical computational abilities, more than normal computer microprocessors. DSPs need to have heavy mathematical computation skills because manipulating analog signals requires it. The second advantage of a DSP lies in the programmability of digital microprocessors. Just as digital microprocessors have operating systems, so DSPs have their very own operating systems. DSPs are used extensively in telecommunications for tasks such as echo cancellation, call progress monitoring, voice processing and for the compression of voice and video signals as well as new telecommunications applications such as wireless LANs and next-generation cellular data and cellular Internet services. They are also used in devices from fetal monitors, to anti-skid brakes, seismic and vibration sensing gadgets, super-sensitive hearing aids, multimedia presentations and desktop fax machines. DSPs are replacing the dedicated chipsets in modems and fax machines with programmable modules — which, from one minute to another, can become a fax machine, a modem, a teleconferencing device, an answering machine, a voice digitizer and device to store voice on a hard disk, to a proprietary electronic phone. DSP chips and DSP cores in custom chips are already doing for the telecom industry what the general purpose microprocessor (e.g. Intel's Pentium) did for the personal computer industry. DSP chips are made by Analog Devices, AT&T, Motorola, NEC and Texas Instruments, among others. DSP cores are made by BOPS, DSP Group, Infineon and others.

Digital Signature A digital signature is the network equivalent of signing a message so that you cannot deny that you sent it and that the recipient knows it must have come from you. In short, a digital signature is an electronic signature which cannot be forged. It verifies that the document originated from the individual whose signature is attached to it and that it has not been altered since it was signed. There are two types of digital signatures. Ones you encrypt yourself and are the result of an ongoing relationship between you and the other party. Second, there are encrypted certificates issued by a com-

MODERN
DICTIONARY
of
ELECTRONICS

RUDOLF F. GRAF

**MODERN
DICTIONARY
of
ELECTRONICS**

SEVENTH EDITION

REVISED AND UPDATED

Rudolf F. Graf




Newnes

Boston Oxford Auckland Johannesburg Melbourne New Delhi


Newnes is an imprint of Butterworth-Heinemann.

Copyright © 1999 by Rudolf F. Graf

 A member of the Reed Elsevier Group.

All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the publisher.

 Recognizing the importance of preserving what has been written, Butterworth-Heinemann prints its books on acid-free paper whenever possible.



Butterworth-Heinemann supports the efforts of American Forests and the Global ReLeaf program in its campaign for the betterment of trees, forests, and our environment.

Library of Congress Cataloging-in-Publication Data

Graf, Rudolf F.

Modern dictionary of electronics / Rudolf F. Graf. — 7th ed.,
revised and updated.

p. cm.

ISBN 0-7506-9866-7 (alk. paper)

1. Electronics — Dictionaries. I. Title

TK7804.G67 1999

621.381'03 — dc21

99-17889

CIP

British Library Cataloguing-in-Publication Data

A catalogue record for this book is available from the British Library.

The publisher offers special discounts on bulk orders of this book.

For information, please contact:

Manager of Special Sales

Butterworth-Heinemann

225 Wildwood Avenue

Woburn, MA 01801-2041

Tel: 781-904-2500

Fax: 781-904-2620

For information on all Butterworth-Heinemann publications available, contact
our World Wide Web home page at: <http://www.bh.com>

10 9 8 7 6 5 4 3 2 1

Typeset by Laser Words, Madras, India

Printed in the United States of America

amplitude permeability — analog network

24

waveshape of random noise and its amplitude is directly proportional to the bandwidth of the transmission system.

amplitude permeability—The relative permeability at a stated value of field strength and understated conditions, the field strength varying periodically with time and no direct magnetic-field component being present.

amplitude range—The ratio, usually expressed in decibels, between the upper and lower limits of program amplitudes that contain all significant energy contributions.

amplitude resonance—The condition that exists when any change in the period or frequency of the periodic agency (but not its amplitude) decreases the amplitude of the oscillation or vibration of the system.

amplitude response—The maximum output amplitude that can be obtained at various points over the frequency range of an instrument operated under rated conditions.

amplitude selection—The process of selecting that portion of a waveform which lies above or below a given value or between two given values.

amplitude separator—A television-receiver circuit that separates the control impulses from the video signal.

amplitude-shift keying—Abbreviated ask. The modulation of digital information on a carrier by changing the amplitude of the carrier.

amplitude-suppression ratio—In frequency modulation, the ratio of the magnitude of the undesired output to the magnitude of the desired output of an FM receiver when the applied signal is simultaneously amplitude and frequency modulated. Generally measured with an applied signal that is amplitude modulated 30 percent at a 400-hertz rate and is frequency modulated 30 percent of the maximum system deviation at a 1000-hertz rate.

amplitude versus frequency distortion—Distortion caused by the nonuniform attenuation or gain of the system, with respect to frequency under specified terminal conditions.

AM rejection ratio—The ratio of the recovered audio output produced by a desired FM signal with specified modulation, amplitude, and frequency to that produced by an AM signal, on the same carrier, with specified modulation index.

AM suppression—The ability of an FM tuner to reject AM signals. Expressed in decibels, it is the ratio between the tuner output with a 100-percent modulation FM signal to its output with a 30-percent modulated AM signal.

AM tuner—A device capable of converting amplitude-modulated signals into low-level audio frequencies.

amu—Abbreviation for atomic mass unit.

analog—1. In electronic computers, a physical system in which the performance of measurements yields information concerning a class of mathematical problems. 2. Of or pertaining to the general class of devices or circuits in which the output varies as a continuous function of the input. 3. The representation of numerical quantities by means of physical variables, e.g., translation, rotation, voltage, resistance; contrasted with *digital*. 4. A continuous representation of phenomena in terms of points along a scale, each point merging imperceptibly into the next. An analog voltage, for example, may take any value. Real-world phenomena, such as heat and pressure, are analog (compare with *digital*).

analog adder—An analog circuit or device that receives two or more inputs and delivers an output that is equal to their sum.

analog amplifier—A device whose output is continuously proportional to the input stimulus.

analog channel—A computer channel in which the transmitted information can have any value between the defined limits of the channel.

analog circuit—A circuit in which the output varies as a continuous function of the input, as contrasted with digital circuits.

analog communications—A system of telecommunications employing a nominally continuous electrical signal that varies in frequency, amplitude, etc., in some direct correlation to nonelectrical information (sound, light, etc.) impressed on a transducer.

analog computer—1. A computer operating on the principle of creating a physical (often electrical) analogy of the mathematical problem to be solved. Variables such as temperature, light, pressure, distance, angle, shaft speed, or flow are represented by the magnitude of a physical phenomenon such as voltage or current. The computer manipulates these variables in accordance with the mathematical formulas “analogued” on it. 2. A computer system in which both the input and output are continuously varying signals. 3. A computing machine that works on the principle of measuring, as distinguished from counting. 4. A computer that solves problems by setting up equivalent electric circuits and making measurements as the variables are changed in accordance with the corresponding physical phenomena. An analog computer gives approximate solutions, whereas a digital computer gives exact solutions. 5. A nondigital computer that manipulates linear (continuous) data to measure the effect of a change in one variable on all other variables in a particular problem. (Compare: *digital computer*.)

analog computing—Computing system in which continuous signals represent mechanical (or other) parameters.

analog data—1. A physical representation of information such that the representation bears an exact relationship to the original information. The electrical signals on a telephone channel are an analog data representation of the original voice. 2. Data represented in a continuous form, as contrasted with digital data represented in a discrete (discontinuous) form. Analog data is usually represented by physical variables, such as voltage, resistance, rotation, etc.

analog input module—An I/O rack module that converts an analog signal from a user device to a digital signal that may be processed by the processor.

analog meter—An indicating instrument that employs a movable coil and pointer arrangement (or equivalent) to display values along a graduated scale.

analog multiplexer—1. Circuit used for time-sharing of analog-to-digital converters between a number of different analog information channels. Consists of a group of analog switches arranged with inputs connected to the individual analog channels and outputs connected in common. 2. Two or more analog switches with separate inputs and a common output, with each gate separately controllable. Multiplexing is performed by sequentially turning on each switch one at a time, switching each individual input to a common output. 3. A device that selects one of several analog signals according to a digital code. Analog multiplexers (amux) are available in many forms; their chief application is as a front end in data-acquisition systems, enabling a single analog-to-digital converter to monitor more than one information channel.

analog network—A circuit or circuits that represent physical variables in such a manner as to permit the expression and solution of mathematical relationships between the variables, or to permit the solution directly by electric or electronic means.

task. Instructions and data are stored in the same memory and both can be manipulated by the computer with equal ease. 8. A device that is capable of solving problems or manipulating data by accepting data, performing prescribed operations (mathematical or logical) on the data, and then delivering or applying the results of these operations.

computer access device input—A device that automatically routes to the computer all teletypewriter observation reports that are received in a standard format.

computer-aided design—See CAD.

computer-aided engineering—Abbreviated CAE. An umbrella term that covers all uses of computers in engineering applications. Thus, computer-aided design and computer-aided manufacturing are branches of computer-aided engineering. The subject area is not usually considered to include software engineering.

computer-aided manufacturing—Abbreviated CAM. The use of computer technology to manage, control, and operate manufacturing either through direct or indirect computer interface with the physical and human resources of the company.

computer-aided software engineering—See CASE.

computer-aided tomography—See CAT.

computer architecture—That set of a computer's attributes (such as registers, addressing modes, and instruction set) that are visible to the programmer.

computer assisted tomography—See CAT.

computer code—Also called machine language. The code by which data is represented within a computer system. An example is binary-coded decimal.

computer control—The parts of a digital computer that have to do with the carrying out of instructions in the proper sequence, the interpretation of each instruction, and the application of signals to the arithmetic unit and other parts in accordance with this interpretation.

computer control counter—1. A counter that stores the next required address. 2. Any counter that provides information to the control unit.

computer diagnosis—The use of data processing systems for evaluation of raw data.

computer entry punch—A combination card reader and keypunch used to enter data directly onto the memory drum of a computer.

computer-generated hologram—A synthetic hologram produced using a computer plotter. The binary structure is formed on a large scale and is then photographically reduced. The holograms are finally etched into a medium.

computer graphics—1. Computer output in the form of pictorial representation (graphs, charts, drawings, etc.) that is displayed visually, usually by a cathode-ray tube. 2. A person-oriented system that uses a computer to create, transform, and display pictorial and symbolic data.

computer-integrated manufacturing—See CIM.

computer interface—1. Peripheral equipment for attaching a computer to scientific or medical instruments. 2. A device designed for data communication between a central computer and another unit such as a PC processor.

computer interfacing—The synchronization of digital data transmission between a computer and one or more external I/O devices.

computerized axial tomograph—See CAT.

computerized robot—A servo model run by a computer. The computer controller does not have to be taught by leading the arm-gripper through a routine: new instructions can be transmitted electronically. The programming for such "smart" robots may include the ability to optimize, or improve, its work-routine instructions.

computer language—1. A system of commands used to develop software for computers (e.g., DOS). 2. The method or technique used to instruct a computer to perform various operations. See high-level language; machine language.

computer-limited—Having to do with the condition in which the time required for computation is greater than the time required to read inputs and write outputs.

computer literacy—1. Computer and information system comprehension. 2. The ability to use computer technology in a particular discipline.

computer network—Two or more connected computers that have the ability to exchange information.

computer numerical control—Abbreviated CNC. The use of a dedicated computer within a numerical-control unit to perform some or all of the basic numerical-control functions.

computer-output microfilm printer—Abbreviated COM printer. A microfilm printer that will take output directly from the computer, thus substituting for line printer or tape output.

computer polarization holography—A technique used to store wavefront information on thin polarization information-recordable materials (i.e., photochromic crystals) by controlling the polarization angle of a small illuminating spotlight in each sampling cell on a crystal.

computer port—The physical location at which the communication line interfaces to the computer.

computer program—A series of instructions or statements prepared in a form acceptable to the computer, the purpose of which is to achieve a certain result. See software.

computer programmer—A person who designs, writes, debugs, and documents computer programs.

computer programming language—A set of precisely defined structures and syntax (representation, conventions, and rules of use and interpretation) devised to simplify communication with a computer, such as BASIC, FORTRAN, C++, and Java. The greater the power of a higher-level language, the greater is the complexity of information that can be precisely conveyed in an efficient manner.

computer science—1. The field of knowledge that involves the design and use of computer equipment, including software development. 2. The science of solving problems with computers.

computer system—The computer and its attached peripherals, such as disk drives, monitor, keyboard, and printer.

computer tape—A high-quality magnetic digital recording tape that must be rated at 1600 fci (flux changes per inch) or 530 flux changes per centimeter, or greater.

computer terminal—Peripheral computer equipment for entering and retrieving data. Sometimes incorporates cathode-ray tube for display.

computer user tape system—See CUTS.

computer utility—A network of central computers linked through data communications facilities to remote terminal systems.

computer word—A sequence of bits or characters that is treated as a unit and that can be stored in one computer location. Same as machine word.

computing—Performing basic and more involved mathematical processes of comparing, adding, subtracting, multiplying, dividing, integrating, etc.

computing device—Any electronic device or system that generates and uses timing signals or pulses of more than 10,000 pulses (cycles) per second and uses digital techniques; inclusive of telephone equipment that uses digital techniques or any device or system that generates

DASD — Abbreviation for direct-access storage device. Any storage device utilizing addressing to let users enter or retrieve data without reference to their physical location. For example, a RAM.

dash — Term used in radiotelegraphy. It consists of three units of a sustained transmitted signal followed by one unit during which no signal is transmitted.

dashpot — 1. A device using a gas or liquid to absorb energy from or retard the movement of the moving parts of a circuit breaker or other electrical or mechanical device. 2. A cylinder and piston device using gas or a liquid to retard the movement of a relay or circuit breaker.

DAT — See digital audio tape.

data — 1. A general term used to denote any or all numbers, letters, symbols, or facts that refer to or describe an object, idea, condition, situation, or other factors. It connotes basic elements of information that can be processed or produced by a computer. Sometimes *data* is considered to be expressible only in numerical form, but *information* is not so limited. 2. A general term for any type of information. 3. Inputs in the form of a character string that may have significance beyond their numerical meaning. 4. Any representations, such as characters or analog quantities, to which meaning might be assigned.

data access arrangement — A protective connecting arrangement that serves as an interface between a customer-provided modem and the switched network. See DAA.

data acquisition — 1. The process by which events in the real world are translated to machine-readable signals. The term usually refers to automated systems in which sensors of one type or another are attached to machinery. 2. The simultaneous collection of data from external sensors, usually in analog form. 3. The function of obtaining data from sources external to a computer system, converting it to binary form, and processing it.

data acquisition and control systems — Assemblies of electronic and mechanical components used to monitor and control complex processes. These systems include the following:

- Process sensors that measure such parameters as temperature, pressure, voltage, and current
- Transmitters that convert measurement data to electrical or pneumatic signals and controls
- Digital computers that test set points, program sequential events, and perform calculations
- Software that provides the computer with instructions and routines
- Process actuators, such as solenoids, relays, valves, and motors, that modify the process in response to computer-generated commands
- Process interface devices, such as analog-to-digital converters, that link transmitters and actuators with digital computers
- Human interface devices, such as printers, keyboards, CRT terminals, switches, alphanumeric displays, chart recorders, and alarms, that facilitate human intervention

data acquisition and conversion system — A method of processing analog signals and converting them into digital form for subsequent processing or analysis by computer or for data transmission.

data acquisition system — 1. A system in which a computer at a central computing facility gathers data from multiple remote locations. 2. System for recording data, usually in digital form, from several sources; can include computing functions.

data bank — A comprehensive collection of libraries of data. For example, one line of an invoice may form an item, a complete invoice may form a record, a complete

set of such records may form a file, the collection of inventory control files may form a library, and the libraries used by an organization are known as its data bank. Synonymous with database.

database — Also data base. 1. The entire body of data that has to do with one or more related subjects. Typically, it consists of a collection of data files (such as a company's complete personnel records concerning payroll, job history, accrued vacation time, etc.) stored in a computer system so that they are readily available. 2. A block of computer memory containing information about one given thing. 3. The collection of current variable data elements defined and maintained by the user. 4. A collection of data, consisting of at least one file, that is sufficient for a given purpose or for a given data-processing system. 5. A large and complete collection of information that covers a variety of subject areas. For instance, a medical diagnostic database might contain symptoms for all common diseases or injuries. 6. A collection of data fundamental to a system or to an enterprise. Made up of comprehensive files of information having predetermined structure and organization and able to be communicated, interpreted, or processed by humans or by automatic means. 7. A collection of related data that can be retrieved from memory at will, such as a mailing list or a list of accounts.

database management — 1. A systematic approach to the storage, updating, and retrieval of information stored as data items, usually in the form of records in a file, where many users, or even many remote installations, will use common data banks. 2. A program that enables a computer to store large amounts of information and then sort it in almost any manner. For example, a company's database could give a list of customers by ZIP code, by credit line, alphabetically by name, or by telephone number. The program takes care of managing the storage and retrieval of the data.

database management system — Abbreviated DBMS. A group of programs that allow users to store, alter, and retrieve information from a database.

database relations — Linkages within a database that logically bind two or more elements in the database. For example, a nodal line (interconnect) is related to its terminal connection nodes (pins) because they all belong to the same electrical net.

data block — Typically, all the data for one item that is entered into a computer for processing, or the computer output that results from processing. An example of an input data block is an individual shipping list; an example of an output data block is a check to be sent.

data break — A facility that permits input/output transfers to take place on a cycle-stealing basis without disturbing execution of the program by a computer.

data bus — 1. A wire or group of wires used to carry data to or from a number of different locations. 2. The output pins of the MPU chip and associated circuitry used for the transmission of data from one point in the system to another. 3. In fiber optics, an optical waveguide used as a common trunk line to which a number of terminals can be interconnected through optical couplers. 4. A system incorporated into fiber-optic communications characterized by several spatially distributed terminals that are served with the same multiplexed signal.

data catalog — A software tool used to list all of the data elements in a database.

data channel (or communication) equipment — Abbreviated DCE. Equipment that interfaces a transmission facility to a transmitting/receiving device. A modem is a DCE.

data code — A structured set of characters used to stand for the data items of a data element, for example,

processes of addition, subtraction, multiplication, and division. 3. A computer system in which circuit operation is based on specific signal levels. In a binary digital computer, there are two such signal levels, one at or near zero and the other at a defined voltage. 4. A device that performs sequences of arithmetic and logic operations on discrete data. 5. A type of data-processing equipment that counts, utilizing numbers to express the values and quantities. General-purpose digital computers include central storage units and peripheral control units and are designed to solve a wide class of problems. A common feature of general-purpose equipment is the ability to externally modify the program of instructions. Special-purpose digital computers are not intended for a typical commercial physical environment and include rugged computers for military and space applications. An analog computer measures cost or conditions. Hybrid computers utilize both modes. 6. A computer that solves problems by operating on discrete representing variables by performing arithmetic and logic processes on this data.

digital data — 1. Data represented in discrete, discontinuous form, as contrasted with analog data represented in continuous form. Digital data is usually represented by means of coded characters (e.g., numbers, signs, symbols, etc.). 2. Any data that is expressed in digits. The term usually implies the use of binary digits.

digital data-handling system — The electronic equipment that receives digital data, operates on it in a suitable manner, records it in a suitable manner on a suitable medium, and presents it directly to a computer or a display.

digital delay line — See active delay line.

digital delay module — See active delay line.

digital delay unit — See active delay line.

digital device — 1. Typically, an IC that switches between two exclusive states or levels, usually represented by logical 1 or 0. 2. An electronic device that processes electrical signals that have only two states, such as on or off, high or low voltages, or positive or negative voltages. In electronics, *digital* normally means binary or two-state.

digital differential analyzer — A special-purpose digital computer that performs integration and that can be programmed for the solution of differential equations in a manner similar to that of an analog computer.

digital disc recording — An analog disc recording that has been made from a master tape recording that was digitally recorded.

digital filter — 1. A linear computation or algorithm performed on a selected series in the form of an input signal that produces a new series as output. The computational device may be a specifically designed electronic system or a conventional computer. 2. Network that operates on discrete samples of a signal to achieve a desired transfer-function operation on that signal. Digital filters divide into two classes: nonrecursive filters produce an output that is a function of only the previous and present inputs; recursive filters produce an output that is a function of both the past and present inputs and outputs.

digital filtering — 1. A computational process or algorithm by which a sampled signal or sequence of numbers, acting as input, is transformed into a second sequence of numbers called the output. The computational process may correspond to high-pass, low-pass, bandpass, or bandstop filtering, integration, differentiation, or something else. The second sequence can be used for further processing, as in a fast-Fourier-transform analyzer, or it can be converted to an analog signal, producing a filtered version of the original analog signal. 2. The process of smoothing, spectrally shaping, or removing noise from a signal. Digital filters are basically mathematical functions

that are performed on the digital data stream; their characteristics can be altered under software control, which adds to their overall flexibility. Finite impulse response (FIR) and infinite impulse response (IIR) are examples of digital filter functions.

digital frequency monitor — A special-purpose digital counter that permits a train of pulses to pass through a gate for a predetermined time interval, counts them, and indicates the number counted.

digital harmonic generation — Abbreviated DHG. The use of circuit elements whose outputs are discontinuous functions of their inputs to produce signals that are an integral multiple of the (fundamental) input signal.

digital image analysis — Technology to measure and standardize the output of a computer-interfaced vidicon system.

digital imaging — The process by which an image that is in electronic form (e.g., a bit-mapped graphic) is altered.

digital information display — The presentation of digital information in tabular form on the face of a digital information display tube.

digital integrated circuit — 1. A switching-type integrated circuit. 2. An integrated circuit that processes electrical signals that have only two states, such as on or off, high or low voltages, or positive or negative voltages. In electronics, *digital* normally means binary or two-state. 3. A monolithic group of logic elements. May be small-scale integration (e.g., SSI gates, flip-flops, latches), medium-scale integration (e.g., MSI decoders, adders, counters), or large-scale integration (e.g., LSI memories, microprocessors). 4. A class of integrated circuits that processes digital information (expressed in binary numbers). The processing operations are arithmetic (such as addition, subtraction, multiplication, and division) or logical (in which the circuit senses certain patterns of input binary information and indicates the presence or absence of those patterns by appropriate output binary signals).

digital integrator — Device for summing or totalizing areas under curves that gives numerical readout. See also integrator.

digital logic modules — Circuits that perform basic logic decisions (AND, OR, NOT); used widely for arithmetic and computing functions, flip-flops, half-adders, multivibrators, etc. See also logic system.

digitally programmable oscillator — A voltage-controlled oscillator designed to accept a digital tuning word instead of the usual analog signal. Internal digital-to-analog (d/a) converter circuits transform the digital input to an analog voltage. Tuning-curve linearization is usually accomplished through a digital memory. The frequency speed is primarily limited by the d/a circuits.

digital modulation — A method of transmitting human voice or other analog signals using a binary code (0s and 1s). Digital transmission offers a cleaner signal than analog technology.

digital multimeter — Abbreviated DMM. A test instrument used to measure voltage, current, and resistance. The readout of measured values is shown on a digital display which is typically a liquid crystal display (LCD).

digital optical processing — The scanning of photographs or transparencies of images, either by a vidicon camera or flying-spot scanner, for the conversion of the images to digital form for storage on magnetic tape.

digital output — An output signal that represents the size of a stimulus or input signal in the form of a series of discrete quantities that are coded to represent digits in a system of numerical notation. This type of output is to be distinguished from one that provides a continuous output signal.

EXHIBIT E



US 20030087636A1

(19) **United States**

(12) **Patent Application Publication**

Mazzara et al.

(10) **Pub. No.: US 2003/0087636 A1**

(43) **Pub. Date: May 8, 2003**

(54) **METHOD FOR PROVIDING MULTI-PATH COMMUNICATION FOR A MOBILE VEHICLE**

(75) Inventors: **William E. Mazzara**, Drayton Plains, MI (US); **Mike A. Hichme**, Rochester Hills, MI (US)

Correspondence Address:
General Motors Corporation
Legal Staff, Mail Code 482-C23-B21
300 Renaissance Center
P.O. Box 300
Detroit, MI 48265-3000 (US)

(73) Assignee: **General Motors Corporation.**

(21) Appl. No.: **10/040,049**
(22) Filed: **Nov. 7, 2001**

Publication Classification

(51) **Int. Cl.⁷** **H04Q 7/20**
(52) **U.S. Cl.** **455/426; 455/522**

(57) **ABSTRACT**

A method and system of providing multi-path communication for a mobile vehicle including at least a primary communication device and a secondary communication device is determined in response to a service request. A capability of the primary communication device and the secondary communication device is determined. A service request from one of the primary communication devices and secondary communication devices is based on the capability determination. A service request is initiated from one of the primary communication devices and secondary communication devices.




FIG. 1




FIG. 2





FIG. 3



US 2003/0087636 A1

May 8, 2003

1

METHOD FOR PROVIDING MULTI-PATH COMMUNICATION FOR A MOBILE VEHICLE

FIELD OF THE INVENTION

[0001] This invention relates generally to the communication of a mobile vehicle. In particular, this invention relates to a method for providing communication for a mobile vehicle over a multiple choice of paths.

BACKGROUND OF THE INVENTION

[0002] A rapidly increasing segment of modern vehicles, such as passenger cars, buses, trains, boats and aircraft, are being equipped with integrated wireless communications systems. Integrated wireless communications solutions enables vehicles to have embedded systems with access to mobile services, such as navigation services, cellular phone services, emergency help/assistance, traffic information, directory assistance services, Internet web access for web browsing and email, remote car diagnostics, anti-theft tracking, in-car office, and other analog or digital voice and data communications applications.

[0003] Such embedded communications devices may have multiple communications paths to select based on various conditions, such as type of service needed (analog vs. digital), communications cost (between available service providers), changing coverage areas and service providers while a vehicle is in motion, and available wireless technology in a given area (cellular phone network-based, satellite-based, radio frequency- or RF-based, etc.). However, some integrated wireless vehicle communications solutions rely solely on the ability of the integrated communication of the embedded system to establish a communications link to a wide area network (WAN).

[0004] Recent advances in wireless technologies have lead to widespread use of portable communications devices. Examples of such devices are data capable cellular phones, bi-directional (2-way) pagers and wireless portable data assistants (PDA). Such portable wireless devices could provide a supplementary wireless communications link between the vehicle and the WAN. This would be advantageous in cases where the embedded system is experiencing service interruption, or cannot provide a certain service type available to the portable device.

[0005] It would therefore be desirable to provide a method for determining the preferred communications device used to establish a communications link from the vehicle to the WAN.

SUMMARY OF THE INVENTION

[0006] One aspect of the present invention provides a method for providing multi-path wide area network access for a mobile vehicle. The primary communication device of the vehicle integrated communications system may access a wide area network (WAN). Upon initiation to the WAN, the secondary communication device availability is determined when the primary communication device queries the service provider. The service type of the secondary communication device is determined when the primary communication device queries the system. The viable battery life of the secondary communication device is determined by the power state and power life indications. The most current

received signal strength indication (RSSI) of the secondary communication device is determined when the primary communication device queries the system. If a determination is made by the primary communication device that a more reliable service is available on the secondary communication device, the pending WAN connection request is initiated by the secondary communication device. The success of the WAN connection is determined when the primary communication device queries the call state of the secondary communication device.

[0007] Another aspect of the present invention provides a system for determining a multi-path wide area network access system for a vehicle. The system may include means for making an initial connection request to a WAN, means for determining the availability of the secondary communication device, means for determining the service availability and service type of the secondary communication device, means for determining the battery life viability of the secondary communication device, means for determining the most current RSSI, means for determining reliability of the available service of the secondary communication device, and means for the secondary communication device initiating a connection request to the WAN. The system may also include means for determining the success of the WAN connection based on the call state.

[0008] Another aspect of the present invention provides a computer usable medium including a computer program code for providing multi-path wide area network access for a mobile vehicle. The computer usable medium may include computer program code that determines if a secondary communication device is available, computer program code that determines service type, computer program code that determines battery life viability, computer program code that determines most relative signal strength indication, and computer program code that determines reliability of the secondary communication device. The program code may also include computer program code that initiates a service request from the secondary communication device.

[0009] The program may also include computer program code that computes a calibrated battery life threshold for determination of viable power. The program may also include computer program code that determines the battery life threshold based on the power state and/or the power life. The program may also include computer program code that computes a calibrated RSSI threshold. The program may also include computer program code that checks for received signal strength indication once the battery life exceeds the calibrated threshold.

[0010] The foregoing and other features and advantages of the invention will become further apparent from the following detailed description of the presently preferred embodiment, read in conjunction with the accompanying drawings. The detailed description and drawings are merely illustrative of the invention rather than limiting, the scope of the invention being defined by the appended claims and equivalents thereof.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] FIG. 1 is a schematic diagram of a system for communicating from a vehicle to a service provider in accordance with the present invention;

US 2003/0087636 A1

May 8, 2003

2

[0012] FIG. 2 is a schematic diagram of another embodiment of a system consisting of a portable network access device linked to an embedded device of a vehicle in accordance with the present invention; and

[0013] FIG. 3 is a flow diagram of one embodiment of a method for determining a multi-path wide area network access for a vehicle in accordance with the current invention.

DETAILED DESCRIPTION OF THE PRESENTLY PREFERRED EMBODIMENTS

[0014] FIG. 1 shows one embodiment of a system for communicating from a vehicle to a service provider in accordance with the present invention at **100**. The system **100** may include one or more vehicle clients **110**, one or more carrier systems (**115**, **120**, **125**), one or more communication networks or wide area networks (WAN) **135**, and one or more service providers **130**.

[0015] Vehicle client **110** may be any suitable vehicle. For example, the vehicle may be an automobile or a passenger carrying unit such as a bus, train, boat or aircraft. The vehicle client **110** may be an embedded device that is capable of providing remote services to the vehicle via a wireless communication link with one or more service providers **130**.

[0016] The carrier system (**115**, **120**, **125**) may be any suitable system for transmitting a signal between a vehicle and a service provider via a communications network. In one embodiment of the invention the carrier system may be a wireless carrier system **115** such as a personal communications system (PCS), a global system for mobile communication (GSM), or the like. In another embodiment of the invention, the carrier system may be a link to one or more satellites **120** that is in communication with one or more base satellite dish receivers **125**. In another embodiment of the invention, the carrier system is a link to one or more satellites **120** that relays the signal between one or more other satellites **120** before communicating with one or more base satellite dish receivers **125**. Other examples of carrier systems are radio links such as microwave, citizen's band (CB), dedicated radio systems such as police, military, or any other suitable radio communications link.

[0017] Communications network **135** may be any suitable system for communicating between vehicle client **110** or carrier systems (**115**, **120**, **125**) and a service provider **130**. In one embodiment of the invention the communications network may be a public switched telephone network (PSTN). In another embodiment of the invention, the communications network may be a multiprotocol Internet protocol (IP) network such as Internet, extranet, private network, virtual private network (VPN), or any other wide area network (WAN) capable of carrying voice and/or digital data in either digital and/or analog format. Alternately the communications network may be a hybrid or virtual communication network.

[0018] Service provider **130** may be any remote system that can provide wireless services to the vehicle client, which may include, for example, a public telephone network. In one embodiment of the invention the service provider may provide navigation services to an embedded vehicle navigation system. In another embodiment of the invention the services provider may provide emergency

assistance using a system such as OnStar. Examples of remote services delivered to the vehicle are navigation services, cellular phone services, emergency help/assistance, real-time traffic information, directory assistance services, Internet web access for web browsing and email, music and video, weather and news reporting, real-time stock market updates, remote car diagnostics, anti-theft tracking, in-car office, and other analog and/or digital voice and/or data communications applications. The service provider **130** may be a single service provider or a combination of several service providers. The service provider may be capable of serving multiple vehicle clients simultaneously.

[0019] FIG. 2 shows one embodiment of a system consisting of a portable communications device (PCD) linked to an embedded device of a vehicle capable of communicating with a service provider in accordance with the present invention at **200**. The system **200** may include one or more embedded devices **205**, one or more embedded communication devices **230**, one or more controllers **210**, and one or more links **240** to one or more PCDs **245**.

[0020] Embedded device **205** may be any integrated service system in the vehicle. In one embodiment of the invention the embedded device may be a system, such as OnStar, capable of providing remote services to the vehicle, such as navigation instructions, roadside assistance, emergency assistance, and directory assistance services. The embedded device **205** may include a global positioning system (GPS) receiver capable of providing vehicle positioning information to the embedded system as well as communicating it to the service provider. In another embodiment of the invention the embedded device may be a GPS based navigation system capable of providing visual and/or audio navigation and map services to the driver from the service provider. Other examples of such embedded devices include cellular phone systems, Internet web access for web browsing and email, audio/video systems such as broadcast and/or on-demand audio and video, text-to-speech news systems, anti-theft systems such as LoJack, remote car diagnostics systems, and integrated personal computer equipment. The embedded device may be a single system or an integration of multiple systems.

[0021] The portable communications device **245** may be a portable network access device (PNAD) capable of wireless communication via a carrier system over a communications network to a service provider. In one embodiment of the invention the PNAD may be a digital and/or analog cellular telephone. In another embodiment of the invention the PNAD may be any wireless communication device including, for example, a web-enabled personal digital assistant (PDA) such as the Palm Pilot, with wireless network access capabilities or a web-enabled wireless phone. The PNAD may be any portable device capable of communicating voice, audio, video, and/or digital data in either digital and/or analog format via a wireless carrier system over a communication network with a service provider. The PNAD may be capable of communicating with at least one embedded system via the physical and/or wireless communication link **240**.

[0022] Link **240** may be any physical or wireless communication link between the PNAD and the embedded system in a vehicle. The link may be capable of communicating voice, audio, and/or digital data in either digital

US 2003/0087636 A1

May 8, 2003

3

and/or analog format between the embedded device and the PNAD. The PNAD may also be capable of communicating device status information and other control information with the embedded device. Examples of device status information of the PNAD are battery life and received signal strength indication (RSSI). In one embodiment of the invention, link 240 may be a physical cable between the PNAD and the embedded device. The cable may be a cable capable of conducting electric and/or electromagnetic signals. The cable may also be a fiber optical cable. In another embodiment of the invention, the link is a wireless communication link. The wireless communication link may communicate using radio signals and/or infra-red light, or the like.

[0023] Embedded communications device 230 may be any embedded wireless transceiver or collection of multiple transceivers such as devices 235a, 235b, and 235c that are part of the embedded system capable of wireless communication via a carrier system over a communications network to a service provider. In one embodiment of the invention, the communications device is an embedded analog and/or digital cellular telephone. In another embodiment, the communications device is a satellite communications device. In another embodiment of the invention, the communications device is a RF transceiver, such as a microwave, a citizen's band (CB) radio, a dedicated radio system for police or military communications, or any other suitable radio communications link. The embedded device may be any device that is capable of communicating voice, audio, video, and/or digital data in either digital and/or analog format via a wireless carrier system over a communications network with a service provider. The embedded communications device 230 may be a hybrid of various communications devices and/or a single device capable of establishing different types of communications links, such as over a cellular telephone network and/or over a satellite radio link.

[0024] Controller 210 may be any control module or device of an embedded service system that is capable of executing program logic for determining which communications device to use in order to establish the communications link with the service provider. In one embodiment of the invention, the controller contains a central processor unit (CPU) 215 that is capable of executing a method stored in memory 220 for determining whether the communications link with the service provider should be established by a PNAD or the embedded device. The controller may be capable of querying status information such as battery life and RSSI of any PNAD linked to the system.

[0025] FIG. 3 shows a flow diagram of one embodiment of a method for achieving a multi-path wide area network (WAN) access to a vehicle in accordance with the present invention at 300. The method illustrated in the embodiment of 300 determines if a more reliable service is available on a portable network access device (PNAD) 245 than the service provided on the embedded device 205.

[0026] The embedded device 205 may be the system master and may handle the initiation and termination of communications sessions. The embedded device 205 may determine if more reliable service is available on the PNAD by querying the PNAD for its battery life status and its related signal strength indication (RSSI). The communications attributes of the PNAD 245 may include the type of service, the battery life, and the relative signal strength indication.

[0027] The embedded system may be capable of accessing various types of WANs, like a connection to a digital network or analog data network. Upon initiation of a WAN connection request, an embedded device 205 in a vehicle may query the system to determine if a portable network access device is available (block 305). PNAD availability may depend on one or more PNADs currently linked to the system being activated. If the PNAD is unavailable, in one embodiment, the embedded device 205 may initiate the WAN connection request (block 350).

[0028] The service type of the WAN connection may be determined prior to initiating the communications session, and may be analog cellular service, digital cellular service, or the like. For example, as new services such as satellite and packet data become available, the service type determination in block 310 may be modified to incorporate them. If the service type is determined to be analog or other non-preferred technology, the embedded communication device may initiate the WAN connection request (block 350).

[0029] When the service type is determined to be digital or the preferred technology (block 310), required battery life threshold of the PNAD 245, which provides the service chosen in 310, may be determined (block 315). The battery life threshold may be determined on parameters such as the power state and power life of the PNAD's power source, as well as the type and duration of the WAN connection. In block 320, the PNAD may be queried for its battery life state. If the battery life state of the PNAD is determined to be insufficient (block 325), the connection may be initiated by the embedded device (block 350).

[0030] If the battery life of the PNAD is determined to be sufficient, (block 325), the required received signal strength indication (RSSI) threshold of the PNAD may be determined (block 330). The PNAD may be then queried for its most current RSSI (block 335). If the most current RSSI does not exceed the required RSSI threshold for the WAN connection (block 340), the embedded device may initiate the WAN connection request (block 350). If the most current RSSI of the PNAD is determined to be sufficient (block 340), the PNAD may initiate the WAN connection request (block 345).

[0031] If the WAN connection was initiated by the PNAD (block 345), the PNAD may be queried (block 355) to determine if the WAN connection was successfully established. If the call state is determined to be unsuccessful by the PNAD (block 355), the embedded device may execute a retry strategy to establish the connection, (block 360). The retry strategy may include trying to re-establish the connection on the PNAD, or it may include initiating the request on the embedded device if the service type is available on the embedded device.

[0032] If the WAN connection was initiated by the embedded device (block 350), the system may determine if the WAN connection was successfully established. If the call state is determined to be unsuccessful (block 355), the embedded device may execute a retry strategy to establish the connection, (block 360). The retry strategy may include

US 2003/0087636 A1

May 8, 2003

4

trying to re-establish the connection on the embedded device, or it may include initiating the request on the PNAD if the service type is available on the embedded device.

[0033] If a successful WAN connection was established by the PNAD, (block 355), both data and/or voice may be communicated via the PNAD which may then communicate with the embedded device via a two-way communications link. If a successful WAN connection was established by the embedded device (block 355), both data and/or voice may be communicated via the embedded communication device.

[0034] While the embodiments of the invention disclosed herein are presently considered to be preferred, various changes and modifications can be made without departing from the spirit and scope of the invention. The scope of the invention is indicated in the appended claims, and all changes that come within the meaning and range of equivalents are intended to be embraced therein.

We claim:

1. A method for providing multi-path communication for a mobile vehicle comprising:

receiving a service request;

determining availability of at least one primary communication device and at least one secondary communication device in response to the service request;

determining capability of the primary communication device and the secondary communication device; and

requesting communication from one of the primary communication device and the secondary communication device based on the capability determination.

2. The method of claim 1, further comprising:

initiating a service request from one of the primary communication device and the secondary communication device.

3. The method of claim 1, wherein the capability determination is based on factors selected from the group consisting of battery life viability, relative signal strength indication, service availability, type of service and call state.

4. The method of claim 3, wherein the battery life viability is based on a power state and a power life.

5. The method of claim 3, further comprising:

determining a calibrated threshold for the battery life viability.

6. The method of claim 5, further comprising:

determining the battery life viability if the calibrated threshold is exceeded.

7. The method of claim 3, further comprising:

determining a calibrated threshold for the received signal strength indication.

8. The method of claim 7, further comprising:

determining the received signal strength indication if the calibrated threshold is exceeded.

9. The method of claim 3, wherein the type of service is analog communication, digital communication, satellite communication, and global system for mobile communication.

10. A system for providing multi-path communication for a mobile vehicle comprising:

means for receiving a service request;

means for determining availability of at least one primary communication device and at least one secondary communication device in response to the service request;

means for determining capability of the primary communication device and the secondary communication device; and

means for requesting communication from one of the primary communication device and the secondary communication device based on the capability determination.

11. The system of claim 10, further comprising:

means for initiating a service request from one of the primary communication device and the secondary communication device.

12. The system of claim 10, further comprising:

means for determining a calibrated threshold for the battery life viability.

13. The system of claim 12, further comprising:

means for determining the battery life viability if the calibrated threshold is exceeded.

14. The system of claim 10, further comprising:

means for determining a calibrated threshold for the relative signal strength indication.

15. The system of claim 14, further comprising:

means for determining the relative signal strength indication if the calibrated threshold is exceeded.

16. A computer usable medium including a program for providing multi-path communication for a mobile vehicle comprising:

computer usable code for receiving a service request;

computer usable code for determining availability of at least one primary communication device and at least one secondary communication device in response to the service request;

computer usable code for determining capability of the primary communication device and the secondary communication device; and

computer usable code for requesting communication from one of the primary communication device and the secondary communication device based on the capability determination.

17. The computer usable medium of claim 16, further comprising:

computer program code for initiating a service request from one of the primary communication device and the secondary communication device.

18. The computer usable medium of claim 16, wherein the capability determination is based on factors selected from the group consisting of battery life viability, received signal strength indication, service availability, type of service and call state.

US 2003/0087636 A1

May 8, 2003

19. The computer usable medium of claim 18, wherein the battery life viability is based on a power state and a power life.

20. The computer usable medium of claim 18, further comprising:

computer usable code for determining a calibrated threshold for the battery life viability.

21. The computer usable medium of claim 20, further comprising:

computer usable code for determining the battery life viability if the calibrated threshold is exceeded.

22. The computer usable medium of claim 18, further comprising:

computer usable code for determining a calibrated threshold for the relative signal strength indication.

23. The computer usable medium of claim 22, further comprising:

computer usable code for determining the relative signal strength indication if the calibrated threshold is exceeded.

24. The computer usable medium of claim 18, wherein the type of service is analog communication, digital communication, satellite communication, and global system for mobile communication.

* * * * *

EXHIBIT F

US006829603B1

(12) **United States Patent**
Chai et al.

(10) **Patent No.:** **US 6,829,603 B1**
(45) **Date of Patent:** **Dec. 7, 2004**

(54) **SYSTEM, METHOD AND PROGRAM
PRODUCT FOR INTERACTIVE NATURAL
DIALOG**

(75) Inventors: **Joyce Yue Chai**, Elmsford, NY (US);
Sunil Subramanyam Govindappa,
Pleasantville, NY (US); **Nandakishore**
Kambhatla, Elmsford, NY (US);
Tetsunosuke Fujisaki, Armonk, NY
(US); **Catherine G. Wolf**, Katonah, NY
(US); **Dragomir Radkov Radev**, Ann
Arbor, MI (US); **Yiming Ye**, White
Plains, NY (US); **Wlodek Zadrozny**,
Tarrytown, NY (US)

(73) Assignee: **International Business Machines
Corp.**, Armonk, NY (US)

(*) Notice: Subject to any disclaimer, the term of this
patent is extended or adjusted under 35
U.S.C. 154(b) by 0 days.

(21) Appl. No.: **09/496,615**

(22) Filed: **Feb. 2, 2000**

(51) **Int. Cl.**⁷ **G06F 17/30**

(52) **U.S. Cl.** **707/5; 707/10; 704/9**

(58) **Field of Search** **707/5, 3, 4, 10,**
707/104.1; 705/26; 704/9

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,426,781 A	6/1995	Kaplan et al.	707/4
5,574,908 A	11/1996	Harding et al.	707/1
5,737,592 A	4/1998	Nguyen et al.	707/4
5,748,974 A *	5/1998	Johnson	704/9
5,784,562 A *	7/1998	Diener	709/217
6,094,652 A *	7/2000	Faisal	707/5
6,173,279 B1 *	1/2001	Levin et al.	707/5
6,182,063 B1 *	1/2001	Woods	707/3
6,233,561 B1 *	5/2001	Junqua et al.	704/277
6,246,997 B1 *	6/2001	Cybul et al.	705/27
6,282,534 B1 *	8/2001	Vora	707/3
6,314,420 B1 *	11/2001	Lang et al.	707/3
6,324,513 B1 *	11/2001	Nagai et al.	704/275

6,346,952 B1 *	2/2002	Shtivelman	345/758
6,377,913 B1 *	4/2002	Coffman et al.	704/8
6,430,531 B1 *	8/2002	Polish	704/257
6,604,141 B1 *	8/2003	Ventura	709/227
6,625,595 B1 *	9/2003	Anderson et al.	707/3

OTHER PUBLICATIONS

Charles Rich et al., "Collagen: A Collaboration Manager for Software Interface Agents," Mitsubishi Electric Information Technology Center America, Mar. 1998, pp. 1-36.

Ion Androutsopoulos et al., "Time, Tense and Aspect in Natural Language Database Interfaces," pp. 1-61.

(List continued on next page.)

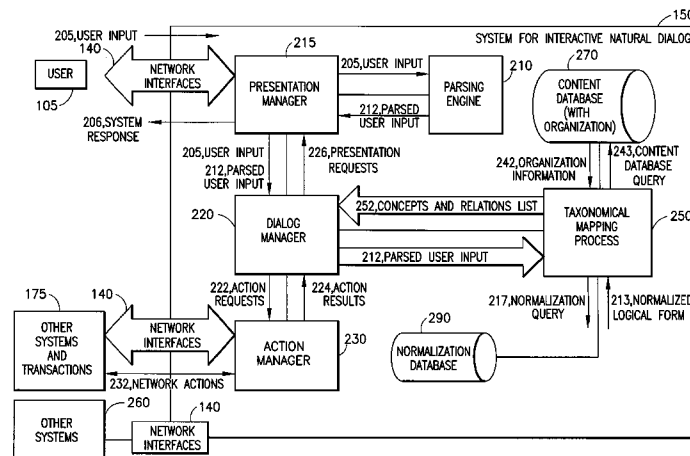
Primary Examiner—Greta Robinson

(74) *Attorney, Agent, or Firm*—Law Office of Charles W. Peterson, Jr.; Louis J. Percello

(57) **ABSTRACT**

This patent describes a novel system, method, and program product that are used in interactive natural language dialog. One or more presentation managers operating on a computer system present information from the computer system to one or more users over network interface(s) and accept queries from the users using one or more known input/output modalities (e.g. Speech, typed in text, pointing devices, etc.). A natural language parser parses one or more natural language phrases received over one or more of the network interfaces by one or more of the presentation managers into one or more logical forms (parsed user input), each logical form having a grammatical and structural organization. A dialog manager module maintains and directs interactive sessions between each of the users and the computer system. The dialog manager receives logical forms from one or more of the presentation managers and sends these to a taxonomical mapping process which matches the items of interest to the user against the content organization in the content database to match business categories and sends modified logical forms back to the dialog manager.

29 Claims, 18 Drawing Sheets



US 6,829,603 B1

Page 2

OTHER PUBLICATIONS

I. Androutsopoulos et al., "Natural Language Interfaces to Databases—An Introduction," Cambridge University Press, 1995, pp. 29–81.

IBM Electronic Commerce, "IBM Net.Commerce Solutions for e-business," Dec. 1997, pp. 1–7.

Samantha Shurety, "IBM e-business with Net.Commerce," (book) Prentice Hall, 1999, pp. 1–718.

[Http://macys.com](http://macys.com), 1 page.

Dragomir R. Radev, "Generating Natural Language Summaries from Multiple On-Line Sources: Language Reuse and Regeneration," Columbia University, 1999, pp. 1–209.

* cited by examiner




FIG.1




FIG.2

```

<PRESENTATION_REQUEST>
  <CLARIFICATION_REQUEST CURRENT_ACTION="INFO">
    <CURRENT_PARAMETERS>
      <PRODUCT GENDER="man" SPECIFIED_BY="user">shirt</shirt</PRODUCT>
    </CURRENT_PARAMETERS>
    <CURRENT_CHOICES>
      <CHOICE>long sleeves</CHOICE>
      <CHOICE>short sleeves</CHOICE>
      <CHOICE>featured Polo shirt</CHOICE>
      <CHOICE>100\% Cotton shirts</CHOICE>
      <NONE_OF_ABOVE>if these items are not what you are looking for, please
        clarify your question or ask a new question</NONE_OF_ABOVE>
    </CURRENT_CHOICES>
  </CLARIFICATION_REQUEST>
</PRESENTATION_REQUEST>

```

FIG.2a




FIG. 2AA

```

<PRESENTATION_REQUEST>
  <QUERY_RESULTS CURRENT_ACTION="INFO">
    <CURRENT_PARAMETERS>
      <PRODUCT GENDER="man" SPECIFIED_BY="user">shirt</shirt</PRODUCT>
      <PRODUCT_TYPE>long sleeves</PRODUCT_TYPE>
    </CURRENT_PARAMETERS>
  </RESULTS>
  <RESULT><PRODUCT>Product 1</PRODUCT></CHOICE>
  <RESULT><PRODUCT>Product 2</PRODUCT></CHOICE>
  <RESULT><PRODUCT>Product 3</PRODUCT></CHOICE>
  <RESULT><PRODUCT>Product 4</PRODUCT></CHOICE>
  <NONE_OF_ABOVE>if these items are not what you are looking for, please
    clarify your question or ask a new question</NONE_OF_ABOVE>
  </RESULTS>
</QUERY_RESULTS>
</PRESENTATION_REQUEST>


```

FIG.2b

INDEX	CONTENT
0	PRODUCT
01	CLOTHING
02	JEWELRY
011	MEN
012	WOMEN
013	KIDS
021	RINGS
022	NECKTIES
023	WATCH
0111	PANTS
0112	T-SHIRTS
0112	SUITS
0114	HATS
0121	PANTS
0122	KHAKIS
0123	SKIRTS
0131	JEANS
0132	HATS

INDEX	CONTENT
0	PRODUCT
01	ACCESSORIES
02	WORKSTATION
03	PERSONAL COMPUTER
011	MONITORS
012	KEYBOARD
021	SYSTEM 8000
031	APTIVA
032	ThinkPad
0321	ThinkPad 770
0322	ThinkPad 600
0323	ThinkPad 570
0324	ThinkPad 390

FIG. 3




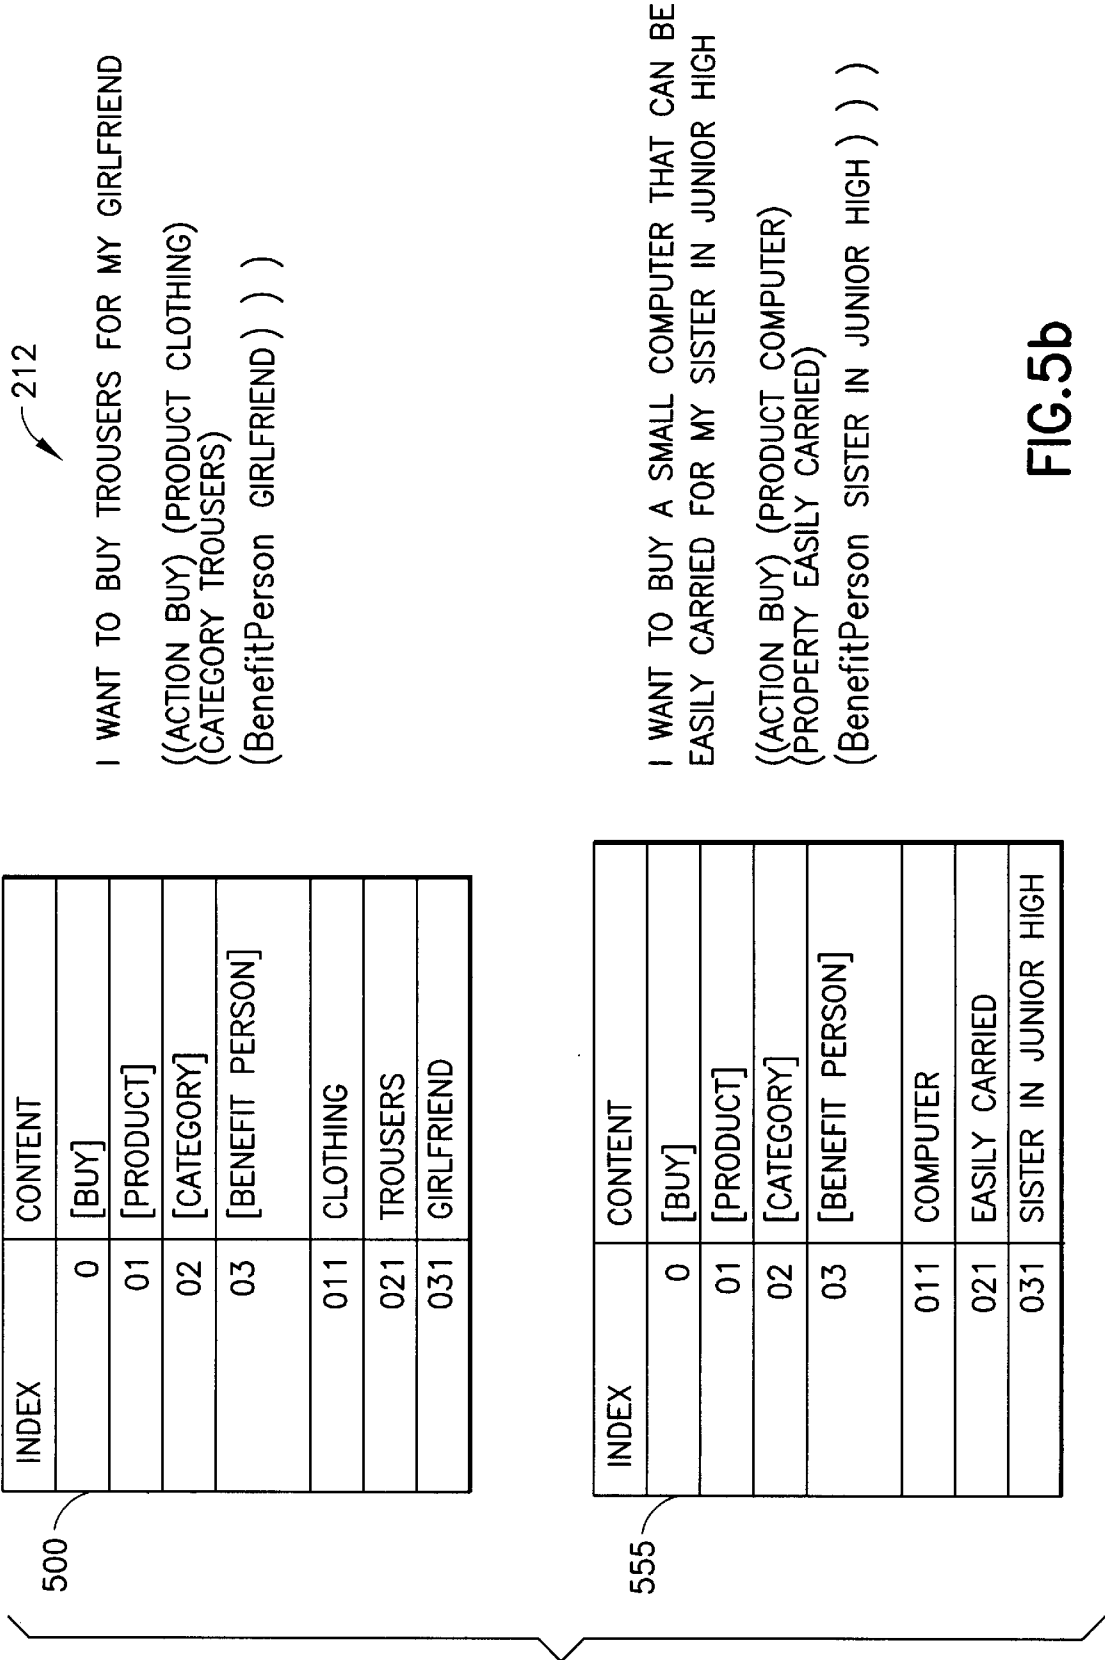


FIG. 5a



213

INDEX	CONTENT
0	[BUY]
01	[PRODUCT]
02	[CATEGORY]
03	[BENEFIT PERSON]
011	CLOTHING
021	PANTS
031	WOMEN

501

INDEX	CONTENT
0	[BUY]
01	[PRODUCT]
02	[CATEGORY]
03	[BENEFIT PERSON]
011	COMPUTER
021	ThinkPad
031	NONPROFESSIONAL

556

FIG. 5c





FIG. 6



U.S. Patent

Dec. 7, 2004

Sheet 13 of 18

US 6,829,603 B1

TRANSACTION:buy
TEMPLATE PARAMETERS:
item name:
sku:
quantity:
size:
color:
Company:
REQUIRED PARAMETERS: item name, quantity, model

FIG.8

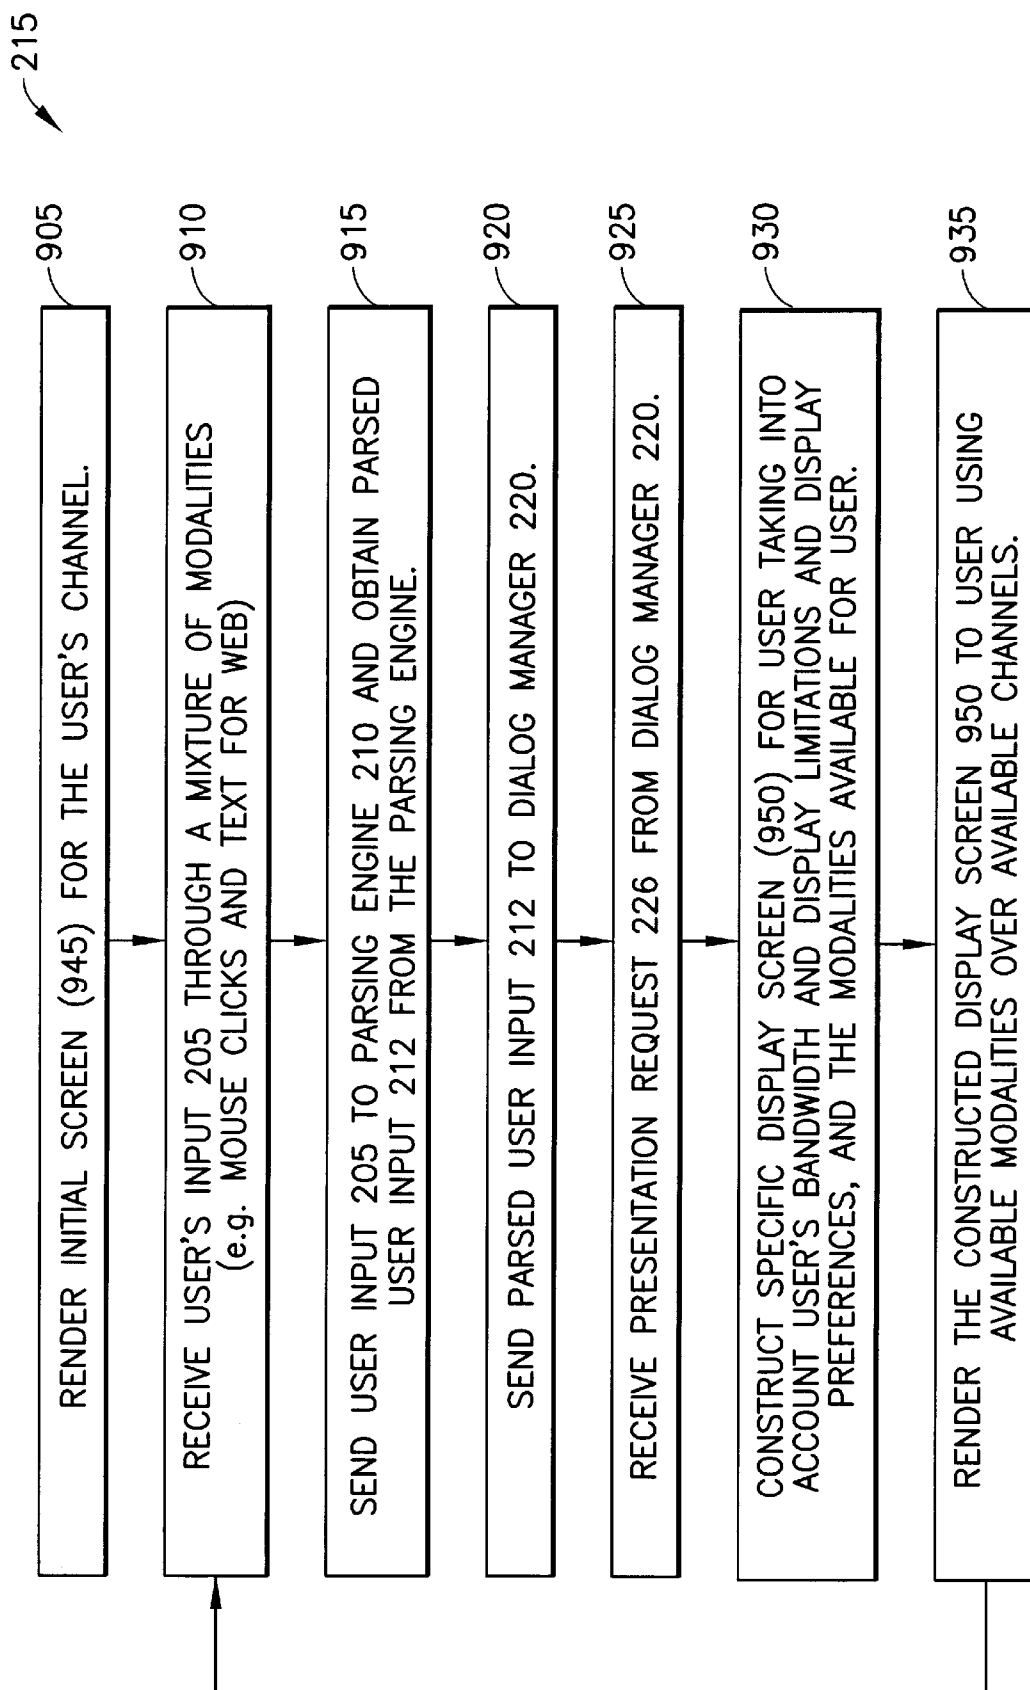


FIG.9

I'm looking for:

(example: a blue striped shirt)

Clear Submit

FIG.9a

WHAT KING OF SHIRTS ARE YOU LOOKING FOR:

- Long Sleeves
- Short Sleeves
- Polo Shirt
- 100% Cotton

IF NONE OF THE ABOVE DESCRIBES YOUR NEED, PLEASE TYPE
IN YOUR REQUEST:

Clear Submit

FIG.9b

WE HAVE FOUND THE FOLLOWING PRODUCTS FOR YOU:

Product 1

Product 2

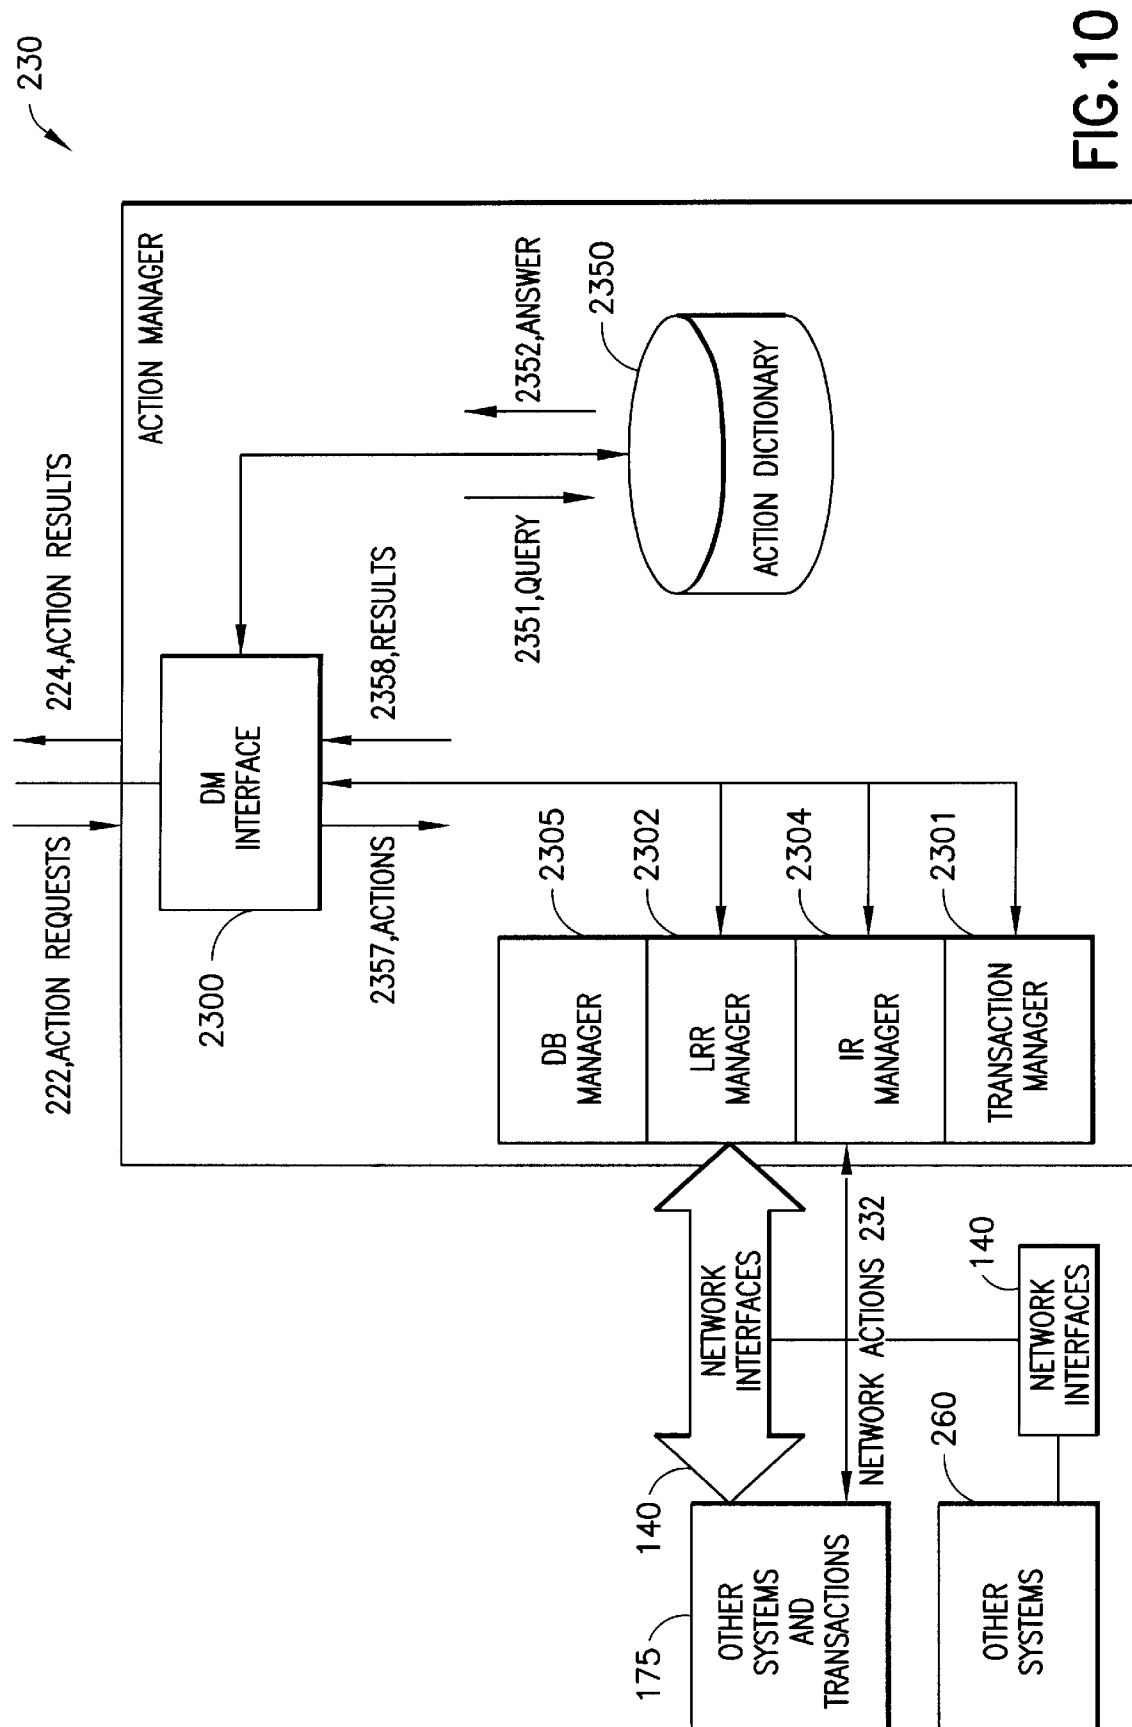
Product 3

Product 4

IF YOU CHANGE YOUR MIND OR LOOK FOR MORE PRODUCTS,
PLEASE TYPE IN YOUR REQUEST:

Clear Submit

FIG.9c



US 6,829,603 B1

1

SYSTEM, METHOD AND PROGRAM PRODUCT FOR INTERACTIVE NATURAL DIALOG

FIELD OF THE INVENTION

This invention relates to database searching and queries, and more particularly to natural language based interactive database searching and queries in network environment.

BACKGROUND OF THE INVENTION

Databases and database search techniques are very well known in the computer arts. Databases have various structures and include any given type of information. In many cases some or all of this information is retrieved by using one or more queries. A query is a request for information from the database that has a structure compatible with the database. Generally, the query is processed in a search that returns results to user.

One common technique for natural language access to databases is to convert natural language sentences to SQL statements. Some examples of SQL statements are shown below:

Query: Show me the names and batting averages of all players who batted above 0.250.

SELECT

 Name, Average

FROM

 Player

WHERE

 Average>0.250

Query: Show me the names and batting averages of all Oriole, Red Sox, and Expo players who batted above 0.300.

SELECT

 Player.Name, Average

FROM

 Player,Team

WHERE

 Average>0.300

AND

 Player.Team=Team.Team

AND

 Team.Name IN ('Orioles', 'Red Sox', 'Expos')

Query: Show me the sum of all batting averages of all players except these from the White Sox and Diamondbacks.

SELECT

 SUM(Average)

FROM

 Player,Team

WHERE

 Player.Team=Team.Team

AND

 Team.Name

NOT IN

 ('White Sox','Diamondbacks')

A paper titled "Natural Language interfaces to databases—an introduction" by I. Androutsopoulos and G. D. Ritchie, appeared in Natural Language Engineering 1(1): 29–81; 1995 Cambridge University Press, which is herein incorporated by reference in its entirety, presents a history of natural language access to databases and provides a survey of the most significant problems that a program that provides such access must face. State-of-the-art database searching includes interactive search, natural language queries and

2

search via internet. One non-natural language interactive database searching technique is described in U.S. Pat. No. 5,426,781 entitled "Computerized report-based interactive database query interface" that discloses a method and system for interactively and iteratively constructing a query using a table metaphor displayed on a user display. Alterations are made directly to the table metaphor by the database user. The alterations relate to adding, deleting, or combining columns of attributes and limiting ranges of attribute values. The alterations are registered and the table metaphor updated to reflect the registered alterations. The table metaphor can be repeatedly used to further register additional alterations. The query corresponding to the table metaphor in its final form is run against the full database to generate a report in the format indicated by the table metaphor.

Using natural language queries to access the information system is also well known. U.S. Pat. No. 5,574,908 entitled "Method and apparatus for generating a query to an information system specified using natural language-like constructs" (herein incorporated by reference in its entirety) discloses an apparatus for generating a query to an information system using a drag-and-drop information system specification means utilizing a computer language having both textual and graphical forms for translating natural language-like constructs into object-role modeling symbol-ogy.

Doing database searching over a general network, e.g. the internet, an intranet, etc. is also well known. In this type of database searching, one or more clients generate a query that is transmitted over the network, a process running on a search processes the query against one or more databases, and returns result to the client back over the network.

U.S. Pat. No. 5,737,592 entitled "Accessing a relational database over the Internet using macro language files" (herein incorporated by reference in its entirety) discloses a method for executing Structured Query Language (SQL) queries in a computer-implemented relational database management system via a network.

One popular way of searching over a network (Internet) is to use a search engine. Most search engines are keyword based search such as YAHOO (<http://www.yahoo.com>), LYCOS (<http://www.lycos.com>) etc., where no user interaction is supported. The user is asked to input the keywords that best represent their interests, then the search engine will look for those keywords (and possibly the synonyms of those keywords) against the document collections. Where a match is found in the document, that document will be retrieved and presented to the user. A typical user is forced to manually go through the many "matches" for a query and find the relevant information herself.

Similar procedures are in place for searching for products. The customers either have to go through a possibly long series of clicking the hyperlinks, or use one of the search mechanisms described above.

Recently, some websites (www.AskJeeves.com, www.Neuromedia.com) have started search operations on question-answer mode. Natural language search engines, such as AskJeeves, use a relatively simple technical approach of keywords, and templates to give the user a feeling of a "natural language interface". For example, a query "What is the capital of Kenya?" returns a pointer to several Web sites including one about Kenya where the correct answer is included in the text. However, a question "How long does it take to fly from London to Paris on the Concorde?" produces a set of similar questions to the one asked however none of them is related to the answer—

US 6,829,603 B1

3

example: "Where can I find cheap flights from the UK?". The method used to produce answers seems to consist of a 5-steps: (a) partly parse the query; (b) map to a canned set of questions/question-templates; (c) map canned questions go to existing knowledge bases (AskJeeves points to other people's web sites for the real data/FAQs.); (d) do a meta search on the 5 big search engines (and return their results too); and (e) if there was no match in "b" then record the query for later human analysis. Note that "b" is essentially a person-intensive task—creating the list of key phrases and the canned questions they map to (and then the underlying web pages they map to). Such systems provide a reasonable front end to a large knowledge base/FAQ. They are better than a raw search engine, because they have the human touch of mapping phrases to canned questions/templates (backed up with the search engines).

Other sites, such as Neuromedia (www.neuromedia.com), BigScience (www.bigscience.com), Novator (www.novator.com), PersonalLogic (www.personallogic.com) try to offer more interactivity to the user. By interactivity we mean the capability of a system to jointly define parameters required for mutual understanding in a series of exchanges. These might be some action parameters, such as Amount, Account_to, Account_from for transferring money, or a set of preferences for a computer notebook. These parameters may be established either by user providing information to the system or the system suggesting some or all of them. What is important is that the system remembers current (and possibly previous) user's preferences, and is using this information in an intelligent manner to make the interaction more satisfying for the user. The above sites, offer more interactivity, by extending the question answer mode of operation with contextual history in the interaction.

PROBLEMS WITH PRIOR ART

The prior art systems fail primarily in three areas:

1. Efficiency: many rounds of interaction are needed to accomplish a task. A typical buying request on average takes about 20 mouse clicks.

2. Lack of deeper understanding of queries. Natural language engines such as AskJeeves cannot be used to accomplish transactions, such as buying clothes, because: (a) a keyword search cannot understand that "summer dress" should be looked upon in women's clothing dept. under "dresses" and "dress shirt" most likely in men's under "shirts", and (b) a search for "shirt" can reveal dozens or even hundred items, which is useless for somebody who has a specific style and pattern in mind. In order to have an appropriate answer, a dialogue with the user is required: the system must come back with questions, e.g., about style, color, etc.

3. Search engines do not accommodate business rules, e.g. a prohibition against displaying heap earrings with more expensive ones.

Search engines, such as AskJeeves, do not engage in dialog with the user, rather they respond to a single question with a set of possible answers. While these search engines are appropriate for searches and can be used in self service in many cases (e.g. For finding general information about the offerings of a bank), in order to have an exact answer—a dialog with the user is required and not sufficient in these systems. For example, to search for a "shirt", sometimes, the user is aware of the specific style and can specify his/her exact interest right at the beginning in a sentence or two; sometimes, the user fails to input the crucial information for

4

the specific style even if he/she has that in mind, sometimes, the user needs help to formulate his/her specific interest. Without dialog with the user to find out his/her exact intention and interest, the search result can reveal dozens or even hundreds of items which could be overwhelming and useless to the user.

For instance, trying to find a pair of black pants without cuffs might take over 10 minutes of exploration of a typical site (such as www.macys.com) because not all pants are available in blacks, some have cuffs and some do not, and also because it is not clear whether the search should be done under "career", "casual" or "active", if the pair in mind is "business casual".

Novator (www.novator.com), PersonalLogic (www.personallogic.com) try to offer more interactivity to the user, for instance in buying a computer a program could ask a number of questions to help with the configuration and pricing. However, the interaction is still unnatural and often time consuming for a casual user, because often it requires spending a lot of time either inputting information or browsing without any feedback from the system about getting closer to the user's goal. Also, the user is required to use and understand the terminology of the site.

OBJECTS OF THE INVENTION

An object of this invention is an improved system, method, and program product for searching computer network sites.

An object of this invention is an improved system, method, and program product for searching computer network sites with fewer user requests to find a proper response to one or more queries.

An object of this invention is an improved system, method, and program product embodying a natural language dialog system that better analyzes and understands queries.

An object of this invention is an improved system, method, and program product embodying a natural language dialog system that allows for a dialog with a user.

An object of this invention is an improved system, method, and program product for searching computer network sites while accommodating domain rules, e.g. business rules.

SUMMARY OF THE INVENTION

This invention is a computer system, method, and program product that has a content database stored on one or more of its memories. The content database has a content organization that may or may not be part of the database. One or more presentation manager modules present information from the computer system to one or more users over one or more network interfaces and accept queries from one or more of the users using one or more known input/output modalities (e.g. Speech, typed in text, pointing devices, etc.). A natural language parser parses one or more natural language phrases received over one or more of the network interfaces by one or more of the presentation managers into one or more logical forms (parsed user input), each logical form having a grammatical and structural organization. A dialog manager module maintains and directs interactive sessions between each of the users and the computer system. The dialog manager receives logical forms from one or more of the presentation managers and sends these to a taxonomical mapping process which matches the items of interest to the user against the content organization in the content database to match business categories and sends modified

US 6,829,603 B1

5

logical forms back to the dialog manager. In a preferred embodiment, these modified logical forms are matched against a set of predefined action templates, the best matching template is selected, corresponding action requests are sent to an action manager, action results are received from the action manager, and presentation requests are sent to one or more of the presentation managers for presenting the system response to the user. The action manager module receives action requests from the dialog manager, executes the action (e.g. Retrieving the latest price of a stock or all blue pants with cuffs), and sends the action results to the dialog manager.

BRIEF DESCRIPTION OF THE FIGS.

The foregoing and other objects, aspects, and advantages will be better understood from the following non limiting detailed description of preferred embodiments of the invention with reference to the drawings that include the following:

FIG. 1 shows a user interacting with an e-commerce site using an interactive natural dialog system.

FIG. 2 shows a high level block diagram of the system architecture of one preferred embodiment of the invention.

FIG. 2AA is a block diagram showing various locations for business rules and/or logic.

FIG. 2(a) shows an example of a presentation request for generating a clarification screen shown in FIG. 9(b).

FIG. 2(b) shows an example of a presentation request for generating a results screen shown in FIG. 9(c).

FIG. 3 shows examples of two tables in the content database.

FIG. 4 shows examples of two tables in the normalization database.

FIG. 5(a) shows a flowchart of the control flow of the taxonomical mapping process.

FIG. 5(b) shows two examples of tabular representations of the parsed user input (logical form).

FIG. 5(c) shows two examples of the normalized logical form.

FIG. 6 shows two examples of querying the content database with a content query in the taxonomical mapping process (shown in FIG. 5(a)).

FIG. 7 is a flowchart of the control flow of the dialog manager.

FIG. 8 shows an example of an action template for a retail 'buy' transaction.

FIG. 9 is a flowchart of the control flow of the presentation manager.

FIG. 9(a) shows an example of an initial screen generated by the presentation manager for a web channel.

FIG. 9(b) shows an example of a display screen for a clarification dialog generated by the presentation manager for a web channel.

FIG. 9(c) shows an example of a display screen for a results display generated by the presentation manager for a web channel.

FIG. 10 is a block diagram of the data flow for the action manager.

DETAILED DESCRIPTION OF THE INVENTION

The present invention improves dialog interaction between the user and the computer in the domain database

6

query formulation and refinement. More specifically, it shows how to construct efficient and effective systems for querying and transactions based on a mapping, or mappings, between the commonsense domain organization and the business organization of data in the said domain.

Another purpose of this invention is to improve interactions of users with networked computers in the domain of electronic commerce and for the internet.

Dialog is one of the most natural ways of interaction with the user. Some interactions with the user involve the user defining SQL language or filling tables or using tools. Typically, these interactions either require advanced skills and therefore exclude the novice users from effective interactions, or are very highly structured, typically based on the business structure of the domain, and therefore precluding fast and efficient interactions. By carrying natural language dialog with the user, generating the language that both skilled and novice users can understand and respond to, a system based on this invention can be used by a wide variety of people.

The present invention is a system and method for carrying such dialogs. Given a domain and a database, we show how to build a system that would allow a user or users to conduct queries and transactions by creating a taxonomical mapping process that associates the grammatical and structural organization of the domain with the content organization of the data. The purpose of the taxonomical mapping is to provide a mapping from conceptual structures of the user—expressed in natural language or a combination of natural language and other media (e.g. pointing, or other sensory data)—into the business organization of the data. And, vice versa, the mapping allows the system to present the business organization using the conceptual structures of the user.

This process addresses the shortcomings of the prior art as follows:

Since natural language allows the user to directly express his/hers intention, instead of navigating the business organization of the database, the invention makes the process more efficient. The capability to express the desire in language reduces the cognitive effort on the part of the user; the fact that it is accomplished through interaction and not just a single query makes the process effective. It is possible to directly map the user request into the most closely related set of items/categories in the database. This is accomplished in one or a few interactions (a dialog), as opposed to 20 or more mouse clicks involving interaction with the server. The efficiency of natural language as an expressive medium has been confirmed in Wizard of Oz experiments. (See e.g. T. K. Landauer "The trouble with computers", MIT Press 1995 pp.282–283.)

The mapping is concerned with relating user's categories with business categories, which cannot be accomplished by keyword search. For example, 'dress' in 'dress shirt' is a modifier and in 'summer dress' is a category. Keyword searches, by definition, cannot distinguish between the two cases, since they only look for words, e.g. dress, without taking into account the category information, which is derived from the grammatical and structural organization of the query.

The problem of dealing with business categories is also addressed by taxonomical mapping, and cannot be solved by prior art. That's because natural language and business categories do not necessarily coincide. For example, the natural language category "earrings" can correspond to two business categories "plastic earrings" and "gold and platinum earrings". Similarly, multiple natural categories, e.g.

US 6,829,603 B1

7

“frogs”, “tofu”, and “vegetables” can correspond to one business category “vegetarian dishes”. Thus having a mapping between the natural language (i.e., grammatical, structural, and semantical) and business organization is essential in addressing the third problem (i.e. that search engines do not accommodate business rules, e.g. a prohibition against displaying cheap earring with more expensive ones).

FIG. 1 shows a high level diagram of the environment 100 of a preferred embodiment of the invention 150. The system for interactive dialog 150 (the system), residing on a server 120, is connected through network interfaces 140 to a transactional database 185 and to a network, preferably the Internet 145. The users 105 would access the system 150 using a (Internet) browser 110 running on a computer 106. In a typical scenario, the invention would be used by an e-commerce site. The browser 110 accesses the system 150 by initiating a session with the e-commerce site. Such a session is maintained by a network interface 140 connecting to one or more of the following: the Internet 145, an intranet, a local area network, a public service telephone network, a wireless cellular network, a cable network, a satellite communications network or any other private or public digital or analog data network.

FIG. 2 shows a high level block diagram of the interactive dialog system 150 (the system). The system receives input 205 from the user 105 through the network interface 140. Typical user input 105 might include typed or spoken requests in natural language for information about specific items or requests to process transactions; e.g. a typed request “show me all blue trousers without pleats”. The presentation manager 215 in the system 150 is responsible for handling all interactions with the user. The presentation manager sends the user input 205 to a parsing engine 210—for different types of parsing see, G. Gazdar and C. Mellish, Natural Language Processing in Prolog, Addison-Wesley Pub. Co., GT. Britain, 1989. or James Allen “Natural Language Understanding” 1995, Addison-Wesley Pub Co.; ISBN: 0805303340. The choice of parsing method and static and dynamic parameters of parsing, such as types of grammar, depth, etc. can depend on other parameters of the system such as the taxonomical mapping, a business model, an information retrieval performance, an information retrieval confidence, actions of the presentation manager, actions of the dialog manager, and actions of the action manager.

The parsing engine 210 in turn parses the sentence and returns the parsed input (logical form) 212 to the presentation manager. The parsed input has a grammatical and structural organization. The parsed input 212 comprises a semantic interpretation of the user’s request. For example, the request 205 “blue trousers without pleats” can be parsed into the logical form structure 212 ((noun: trousers)(modifiers (color blue)(property (neg (noun pleats)))). After receiving the parsed input 212 from the parsing engine 210, the presentation manager 215 sends the parsed input 212 to the dialog manager 220 for interpreting the user’s input in the context of the present conversation.

The dialog manager 220 maintains the state of the current conversation in its internal memory and is responsible for controlling the entire interaction with the user through the presentation manager 215. Upon receiving any new parsed user input 212 from the presentation manager, the dialog manager 220 sends the parsed user input 212 to a taxonomical mapping process 250 to determine the exact business product category or categories asked for.

The taxonomical mapping process 250 receives the parsed user input 212 from the dialog manager and issues a query

8

243 based on the parsed user input 212 to the content database 270 containing an organization specific tabular mapping between category names and the business categories (for instance “trousers” may be mapped to “pants”). The results of the query 242 are sent back to the taxonomical mapping process 250 which then sends concepts and relations list 252 to the action manager 230. Thus the example parsed output 212 ((noun: trousers)(modifiers (color blue)(property (neg (noun pleats))))) can be converted into 252 ((product: pants)(modifiers (color blue) (negative pleats))).

The dialog manager 220 receives the concepts and relations list 252 from the taxonomical mapping process 250. Using the concepts and relations list 252, the dialog manager updates its internal state information and determines if further clarification is needed from the user before fetching items/information from the database. For instance, for a stock trading application, when the user 105 is requesting to buy a stock, the dialog manager 220 first checks the parsed user input 212 and its internal state information to see if all the parameters of the buy operation have been specified by the user. If some clarification is required from the user, the dialog manager 220 sends a presentation request 226 to the presentation manager 215 to ask the user 105 for clarifications about his/her request. FIG. 2(a) shows an example of a presentation request 226 for a clarification dialog with the user that results in the screen (system response 206) in FIG. 9(b) being shown to the user.

If all the parameters of the request have been identified, the dialog manager 220 sends an action request 222 to the action manager 230 to process the information/transactional request. For instance, the action request 222 might be a message requesting the action manager 230 to execute a stock “buy” transaction (“buy IBM shares 150 limit 104 valid today”) or a request to retrieve all items of the type “pants (color blue) (pleats no)”. The dialog manager then waits to receive action results 224 from the action manager and sends appropriate presentation requests 226 to the presentation manager 215 to present the system response 206 to the user’s 105 original query. FIG. 2(b) shows an example of such a presentation request 226 sent to the presentation manager 215 that contains the system responses to a user’s query. This presentation request 226 results in the results screen (system response 206) shown in FIG. 9(c) being shown to the user.

The action manager 230 is responsible for receiving action requests 222 from the dialog manager 220, processing them (i.e. Execute the transactions or retrieve the information), and returning action results 224 (containing the results of the action) to the dialog manager. An example of an action request 222 (a news request, in a stock buying application) is:

```

<ACTION_REQUEST>
  <USER_INPUT>any news on Cisco today?</USER_INPUT>
  <PARSED_INPUT>((action news)(stock cisco))</
  PARSED_INPUT>
  <REQUEST>
    <COMPANY_NEWS TIME="TODAY"
      SYMBOL="CSCO">
    </COMPANY_NEWS>
  </REQUEST>
</ACTION_REQUEST>

```

The above action request from the dialog manager 220 results in the following action result 224 message being sent to the dialog manager from the action manager 230:

US 6,829,603 B1

9

```

<ACTION_RESULT>
  <COMPANY_NEWS>
    <TIME DAY="5" MONTH="October"
      HOUR="4:48PM"/>
    <COMPANY SYMBOL="CSCO" CHANGE="-7 7/16"
      PRICE="48 5/16" VOLUME="1,200,000">
    <LRR> Shares of Cisco Systems Inc. (CSCO) plummeted
      7-7/16 to close at 48-5/16 after the company
      confirmed that the FTC is investigating the
      company.
    </LRR>
  </COMPANY_NEWS>
</ACTION_RESULT>

```

After receiving an action request from the dialog manager, the action manager formulates network actions **232** (e.g. an SQL query to a transactional database, a query to an information retrieval engine etc.) that are sent over the network interface **140** (e.g. The internet, a LAN, ethernet connection, remote dialup connection, etc.) and obtains network actions **232** in reply that contain the results of the action. The results **252** and **232** are sent by the action manager **230** in the form of action results **224** to the dialog manager **220** for presentation to the user. In the examples of the action request and action result shown above, the action request **222** from the dialog manager **220** is routed by the action manager **230** and to other systems **260**, other transactional systems **175**. The action result **224** that the action manager **230** sends back to the dialog manager **220** is a reformulated version of the output of the other systems (**175**, **260**).

When the user **105** inputs a clarification to his/her request or inputs another request, the process described above is repeated.

As shown in FIG. 2AA, business rules/logic **2150**, **2200**, **2500**, are used in several modules. First, in possibly adding contextual information to user input, e.g. in dialog manager **220**, requesting a more expensive set of merchandise to be displayed to more affluent users. Second, in deciding how data/answers will be presented to the user, in presentation manager **215**, e.g. showing promotional items more prominently. Third, in the taxonomical mapping process **250**, in deciding which responses from the database should be presented, e.g. no mentioning of plastic earrings if the query comes after visiting jewelry web page. Also, business rules apply to query mapping into a concept and relationship list, in **250**, e.g. "cheap computer" means "cheaper than \$1200 on every day except Monday when it means cheaper than \$1000". The above possibilities apply to both final and partial answers, as well as request for elaboration.

FIG. 3 shows two examples of tables stored in the retail content organization **270**: Table **315** and Table **355**, where Table **315** shows the content organization of a computer retailer and Table **355** shows the content organization of a clothing retailer. In the preferred embodiment, the content organization includes any one or more of the following: a taxonomy of a web site (hierarchically structured grouping of Web pages), a business model organization, a taxonomy of products (hierarchically structured grouping of products), a taxonomy of services (hierarchically structured grouping of services), one or more product categories, one or more service categories, a product list, a service list. In the preferred embodiment, the content organization structure is in the form of a tree (however other organizations, e.g. directed acyclic graphs, or associative memories could also

10

be used). This tree is stored in the computer's memory in the form of a table, as in the content database **270** of FIG. 2. In a preferred embodiment, the tables have two columns. The first column represents the positions of nodes in the tree. The index $a_1 \dots a_{n-1}$ gives the path of a node in the tree. a_1 is always 0, which refers to the root of the tree. $a_1 a_2$ refers to the node that is the a_2 'th node from left to right on the second level of the tree. The position of $a_1 \dots a_{n-1}$ can be inferred similarly. The value of the content part gives the value of the node.

In the preferred embodiment, the content database **270** contains one or more of the following: service information, product information, retail information, wholesale information, one or more product images, text, voice or video information.

FIG. 4 shows two examples of tables stored in the normalization database **290**. These tables associate natural language expressions with their normalized representations. This database is used by the taxonomical mapping process **250** to obtain the normalized logical form **213** of the parsed user input **212**. Table **415** shows an example of normalized expressions based on classifying words based on linguistic knowledge (synonymy and hyponymy); table **417** uses normalization based on the knowledge of the domain. The two methods can also be used in combination. All of it is prior art.

FIG. 5(a) shows a high level diagram of the taxonomical mapping process **250**. The taxonomical mapping process **250** receives the parsed user input **212** from the action manager and converts it into a tabular representation. FIG. 5(b) shows examples of tabular representations of parsed user input **212**. Table **500** shows the tabular representation of the parsed user input **212** of the user query **205**: "I want to buy trousers for my girlfriend". The table is a tree of the parsing result. The meaning of the index and the content is exactly the same as the description of Table **315** and Table **355** in FIG. 3. Table **555** shows the tabular representation of the parsed user input **212**, for the natural language query: "I want to buy a small computer that can be easily carried for my sister in junior high" in the computer storage in the form of a table.

Referring to FIG. 5(a), the tabular representation of the parsed user input **212** is used by step **511** to formulate the normalization query **217**. This normalization query can be the same as the parsed user input **212**, or it might specify additional information, e.g. formatting information (prior art). In step **521**, the normalization database **290** is queried using the query **217** which results in the normalized logical form **213** of the normalization query **217** (and therefore of **212**). —The examples of such normalized logical forms for the above two queries are presented as **501** and **556** in FIG. 5(c).

The next step, **531**, uses **217** to formulate a content query **243** against the content database **270**. As above, this query **243** can be the same as **217**, or might specify additional information, e.g. formatting information (prior art). The result of querying **270** in step **541** are presented as organization information **242**, which in turn is used in step **551** to build **252** concepts and relationship list, representing the content and organization of the database **270**. As before, in the preferred embodiment **252** might be the same as **242**.

FIG. 6 describes step **541**. In this step, content queries **243(a)** and **243(b)** are matched using rules **573** and **575** (respectively). These rules describe how structured queries such as **243(a)(b)** should be matched with the structure of content. For example, rule **575** says that a thinkpad of choice

US 6,829,603 B1

11

for a nonprofessional person should be either model **570** or **390**. In the preferred embodiment, these rules are part of content organization **270**. However they could also reside on a separate database. Such rules are prior art. The result of matching the rule (or rules) is the required organization information **242**, i.e. the classification of parsed user input **212** in terms of the business taxonomy **270** (content organization).

After executing steps **541** and **551** as described above, the taxonomical mapping process **250** sends the concepts and relations list **252** to the dialog manager **220** as described earlier.

The dialog manager **220** maintains the context of interaction (the session context) between the each of the users and the computer system; the context comprising one or more of the following: a session memory, a transaction history (history of transactions done by user, like a request to pay a bill in a banking domain), a presentation history (history of way the response has been presented to the user by presentation manager), an abstracted session memory, a discourse model (a module for discourse analysis), dialog planner (a module which decides what the dialog with the user should look like), a list of goals, a list of user intentions, a list of subdialogs, and a business logic governing the interaction (cf. Allen 1995, for definition of all these terms). The dialog manager allows users to recover from errors during the interaction, based on the session context.

FIG. 7 shows a flow chart of the dialog manager process **220**. In step **805**, the dialog manager receives the parsed user input **212** from the presentation manager **215**. The dialog manager then sends the parsed user input **212** to the taxonomical mapping process **250** and receives the concepts and relations list **252** in reply. In step **810**, the concepts and relations list **252** is compared against a fixed set of application specific action templates **806** and the matching templates are chosen. In the preferred embodiment, this matching is performed by using pre-defined rules and a matching criterion against pre-defined templates. In other embodiments, a standard supervised machine learning algorithm e.g. neural networks (well known prior art) can be used to learn the matching rules automatically from a corpus of labeled (enhanced logical form, filled template) pairs.

The dialog manager supports automatic generation of follow-up questions based on one or more of the following: the discourse, presentation history and domain lexicon (terminology used in the domain). The dialog manager also supports two or more users engaged in simultaneous interaction with the system **150**, where the users are in synchronous collaboration (PRIOR ART; as in e.g. "COLLAGEN: A Collaboration Manager for Software Interface Agents", by Charles Rich and Candance L. Sidner, *User Modeling and User-Adapted Interaction, Special Issue on Computational Models for Mixed Initiative Interaction*, March 1998, incorporated here in its entirety) with each other and with the computer system by sharing the same session or sharing parts of their respective sessions.

FIG. 8 shows an action template **806** for buying retail items. The concepts and relations list **252** is matched against this template (and all other application templates **806**) to determine the instantiations of slot **807** values. For example, the template in FIG. 8 is instantiated if the values of slots **807** "item name", "quantity", "model" and "transaction" are instantiated with appropriate values. Thus, a "buy" template is a valid match only if the value of the slot "transaction" is "buy".

Referring to step **815** in FIG. 7, after the matching templates **806** are chosen, the dialog manager **220** process

12

determines the slots **807** whose values remain uninstantiated. The dialog manager performs discourse analysis (as in e.g. Allen 1995), which means looking at the discourse history of the current user session, the history of the user interactions in previous sessions, and the preferences of the user to determine if any of the missing slot **807** values can be inferred from context. e.g. the company name for buying a stock might be inferred from the context of the natural language conversation with the machine. In the above example, the value of the attribute "size" can be inferred from the long term history of interactions with the current user and the value of the attribute "color" may be inferred from the conversational history of current session (e.g. the user might say "show me blue shirts" and "I would like to buy two of these", where in the second sentence, the user really means "I would like to buy two of these blue shirts"). After filling slot values based on context (if possible), in step **820**, the dialog manager chooses the best template **806** from among the matched templates based on some best choice criteria. Examples of such criteria include choosing the template with most slots filled, choosing the template with the most number of required slots filled, and choosing the template closest (semantically) with most recent template. If multiple templates remain even after applying the criteria, the dialog manager sends a presentation request **226** to the presentation manager **215** to ask the user to disambiguate his/her query.

Suppose the user input **205** was "I want to buy one 30 inch wide White Color Cooktop with Electric PowerSource and with Electric:Radiant Glass CookTop Surface and of Ken-More company" followed by another request "I want to buy one Black CookTop of previous type" Before taking discourse context into consideration, the following slots **807** in matching templates **806** with "buy" transaction are filled.

First Template (Used for Pants and Shirts):

TRANSACTION: buy

TEMPLATE SLOTS:

item name: CookTop

sku:

quantity: 1

size:

color: Black

company:

REQUIRED SLOTS: item name, quantity, company

Second Template:

TRANSACTION: buy

TEMPLATE SLOTS:

item name: CookTop Power Source:

Cooktop Surface:

Venting:

sku:

quantity: 1

size:

color: Black

company:

REQUIRED SLOTS: item name, quantity, company

After using the discourse context and filling the remaining slots **807**, the templates will look like this

First Template (Used for Pants and Shirts):

TRANSACTION: buy

TEMPLATE SLOTS (i.e., PARAMETERS):

item name: CookTop

sku:

quantity: 1

size: 30 inch

color: Black

company: Kenmore

US 6,829,603 B1

13

REQUIRED SLOTS: item name, quantity, company

Second Template:

TRANSACTION: buy

TEMPLATE SLOTS:

item name: CookTop

Power Source: Electric

Cooktop Surface: Electric: Radiant Glass

Venting:

sku:

quantity: 1

size: 30 inch

color: Black

company: Kenmore

REQUIRED SLOTS: item name, quantity, company

Both the templates **806** have all the required slots **807** filled after the discourse context is taken into consideration but since the second template has more parameters filled, the second template is chosen as the best matching template in the preferred implementation.

In step **825** of FIG. 7, the chosen best matching template is examined to see if all its required slots have been instantiated. If some of the required slots are missing, a presentation request **226** is sent (Step **830**) to the Presentation Manager **215** to ask the user for the missing information. After all the required slots are filled, the completely instantiated action template **806** is sent (Step **840**) to the Action Manager **230**. The dialog manager **220** then receives an action result message **224** from the action manager (step **850**) and sends the same (step **860**) as a presentation request **226** to the presentation manager **215** for display to the user as system response **206**.

FIG. 9 shows the control flow of the presentation manager **215**. The presentation manager is responsible for obtaining any input from the user **205** and for displaying the system's response **206** to the user. The idea of separating presentation from content is in the prior art. However, the integration of different modalities and channels is new. In step **905**, a welcome display **945** for the user's specific channel is rendered. In the case of Web interaction, a welcome screen **945** is displayed (an example is shown in FIG. 9(a)). In step **910**, the presentation manager **215** receives user's input **205** through one or more modalities of interaction (e.g. Keyboard input, keyboard output, speech input over a telephone, speech output over a telephone, speech input through a microphone, speech output over speakers, mouse input, a pointing device input, a dataglove, a device for translating signals into digital data, etc.). For example, in the Web interaction, users can use different modalities that include mouse clicks, screen touches, text input and so on. In step **915**, the presentation manager pre-processes the user input **205** (e.g. run speech recognition), sends it to the parsing engine **210** and obtains the parsed user input **212** (which is independent of modalities and channels) from the parsing engine. For instance, if the user input is spoken, the presentation manager **215** is responsible for executing a speech recognition process to obtain a textual representation of the user's utterance. Then the presentation manager employs a natural language parser **210** and semantic classification to convert the textual input into parsed user input **212** that is sent to the dialog manager (Step **920**).

The presentation manager **215** is also responsible for obtaining presentation requests **226** from the dialog manager **220** and presenting it to the user **105** as the system response **206** using appropriate channels and one or more modalities of interaction. In step **925** of FIG. 9, the presentation manager receives presentation requests **226** from the dialog manager. In step **930**, the presentation manager constructs a

14

specific display screen **950** by taking into account the bandwidth of the user's display device, limitations of the user's display device, personal preferences, and the modalities available for the user. For instance, if the user had spoken her request, the presentation manager **215** might decide to present the system's response **206** in an audio format by executing a "text-to-speech" process. The choice of the specific output format is based on the different parameters in step **930**. Thus, the presentation manager **215** might display the system response **206** as a HTML table, as a textual description, as a spoken summary, etc. The user preferences might either be inferred by the system or explicitly stated by the user (through some mechanism for specifying preferences). Finally, in step **935**, the presentation manager **215** renders the constructed display screen **950** to the user using available modalities over available channels. After displaying the system response, any user input **205** (e.g. a clarification or a correction or a new request) is again sent to the presentation manager **220** as described above. FIGS. 9(b) and 9(c) show examples of display screens **950** for a web channel for a clarification and results screen respectively.

FIG. 10 shows a block diagram of the action manager **230** module. The Action manager (AM) (**230**) maintains communication and transactions with one or more of the following systems: an information retrieval system, a knowledge base (database) of documents, a relational database, a directory of information (e.g.: A group of categorized URLs), an internet site, or any other computer system. The Action manager (AM) (**230**) is in charge of communication with back-end applications such as database managers. More specifically, it expects to receive through its DM interface (**2300**) an action request **222** from the Dialog Manager and channels it to one or more of its subordinate modules. To that effect, it looks up the type of action (**2357**) of the action request in the action dictionary (**2350**) by sending a query (**2351**) and receiving an answer (**2352**) and decides which module it should be routed to. Subordinate modules to the Action Manager include but are not limited to **2301** (transaction manager), **2302** (language reuse and regeneration (cf Dragomir R. Radev. Language Reuse and Regeneration: Generating Natural Language Summaries from Multiple On-Line Sources. PhD thesis, Department of Computer Science, Columbia University, New York, October 1998. included it in its entirety), **2303** (database manager), and **2304** (information retrieval manager). The subordinate managers (**2301**–**2305**) communicate through one or more network interfaces (**140**) with external systems (**175,260**) and perform a number of network actions (**232**). The action manager **230** is responsible for obtaining the results of the subordinate managers and checking whether the request was successful. If yes, it may perform additional processing on the return message and route it back (**2358**) to the Dialog Manager. If no, it must notify the Dialog Manager **220** of the failure. This communication from the action manager to the dialog manager **220** is in the form of action results **224** messages.

The communication with the back end host is prior art (Client/Server Programming with Java and CORBA, Second Edition by Robert Orfali, Dan Harkey ISBN: 047124578X).

For completeness of the description of the preferred embodiment, we describe a possible implementation.

Operation of the Invention

In its preferred embodiment, the invention supplements Web sites by providing an interactive mechanism (a window or applet) to support natural language interactions with data

US 6,829,603 B1

15

stored in databases on the Web. In its preferred embodiment the interactions include transaction or information requests (including navigation, i.e., getting the user to an appropriate web page or document), or any combination of the above. One main idea of the present invention lies in realization that the key to effective and fast interactions lies in providing support for both customers taxonomies (needed for natural language dialog) and business taxonomies (used e.g. for planning and reporting). The present invention provides such support by creating an association relation (or relations) between the two types of taxonomy. The user can specify an item or items and its parameters in natural language; the system responds showing appropriate pages (dynamically constructed) containing either the answer to the user query/request or a request to provide more information. The invention supports ungrammatical and misspelled queries in the natural language window. (The input can be typed there, but if speech recognition is available, the users could speak their queries). The system responses combine pictures, choice boxes and natural language. Parts of the system response can be given in recorded voice. The system can support personalization of output and input (e.g. sizes, age, color preference). The system can update the dialog automatically, e.g. by introducing sales items, and personalizing them based on the customer's data, the dialog so far, and merchant objectives. The system can support imprecise queries and requests, e.g. "I need a gift for my mom". How many questions can be supported is limited by the amount of descriptive information the merchants put in their database.

Recorded voice segments can be added to enhance the experience and direct customers attention. The following are typical scenarios:

Scenario 1: When the customer has specific items in mind, he/she would like to get the desired items quickly. The virtual agent tries to understand customers' special requests in NL and helps to find items efficiently without going through navigation. In addition, the virtual agent is able to answer specific questions customers might have in order to make decisions.

Scenario 2: When the customer only has some ideas but not quite specific, he/she would like to get some suggestions from the agent and also see more items to make a choice. The virtual agent figures out the customer's interests by asking questions related to the features of merchandises. Based on the responses, the virtual agent finds the items or information about classes of items.

Scenario 3: When the customer has multiple items in mind but not quite specific, he/she might want to see how those items match. The virtual agent first shows the matching items and then finds out customers' desired items by conversation.

Scenario 4: When the customer has a very general idea, such as "gift", he/she probably likes to get help from the virtual agent. The virtual agent takes the initiative by asking customers questions and lead customers to some category.

We elaborate Scenario 1: Find Request

User: I'm looking for a blue striped shirt and navy pants.

System: presents radio button selections for: men's, women's, boys', girls'.

User: selects men's

System: based on past buying history, presents a selection of Calvin Klein shirts and pants.

This screen allows user to select an item to see more details, buy it, or ask another question.

User: selects a pair of pants.

16

System: presents details for that pair of pants, box to buy it, and box to ask another question.

User: Do you have these in a lighter shade?

System: presents screen with additional pants. This screen allows user to select an item to see more details, buy it, or ask another question.

User: user selects a pair of pants and clicks to buy it and fills in size details.

System: screen shows that pants have been purchased. Shows boxes for: "show me the shirts again", and "show me some other shirts", "show me the matching jacket", and box to ask another question.

User: chooses matching jacket.

System: presents matching jacket details, box to buy it, and box to ask another question.

User: clicks to buy and fills in size details.

System: screen shows that jacket and pants have been purchased.

While the invention has been described in terms of preferred embodiments, those skilled in the art will recognize that the invention can be practiced with modification within the spirit and scope of the appended claims.

We claim:

1. A computer system with one or more memories, one or more central processing units (CPU), and one or more network interfaces, the system further comprising:

a content database stored on one or more of the memories, the content database having a content organization;

one or more presentation manager modules that present information from the computer system to one or more users over one or more of the network interfaces and accept queries from one or more of the users using one or more modalities over one or more of the network interfaces; a

a parsing engine that parses one or more natural language phrases received over one or more of the network interfaces by one or more of the presentation managers, the natural language phrases being parsed into one or more logical forms, each logical form having a grammatical and structural organization;

a dialog manager module that maintains and directs interactive sessions between each of the users and the computer system by iteratively:

receiving logical forms from one or more of the presentation managers,

performing a discourse analysis on these logical forms, selectively generating follow-up questions, said follow-up questions being presented to a respective user after each iteration until said dialog manager determines further clarification is unneeded,

sending action requests to an action manager responsive to a determination that further clarification is unneeded,

receiving action results from the action manager, and sending presentation requests to one or more of the presentation managers;

an action manager module that receives action requests from the dialog manager, performs the action, and sends the action results to the dialog manager, and

a taxonomical mapping process that associates the logical forms of the natural language phrases to the content organization of the content database.

2. A computer system, as in claim 1, where the network interface connects to one or more of the following: the Internet, an intranet, a local area network, a public service telephone network, a wireless cellular network, a cable

US 6,829,603 B1

17

network, a satellite communications network, one or more other private digital data networks, one or more other public digital data networks, one or more other private analog data networks, and one or more other public analog data networks.

3. A computer system, as in claim 1, where the content organization includes any one or more of the following: a taxonomy of a web site, a business model organization, a taxonomy of products, a taxonomy of services, one or more product categories, one or more service categories, a product list, and a service list.

4. A computer system, as in claim 1, where the content database contains one or more of the following: service information, product information, retail information, whole-sale information, one or more product images, text, voice information, and video information.

5. A computer system, as in claim 1, where the dialog manager maintains a session context between each of the users and the computer system, the session context comprising one or more of the following: a session memory, a transaction history, a presentation history, an abstracted session memory, a discourse model, dialog planner, a list of goals, a list of user intentions, a list of subdialogs, and a business logic governing the interaction.

6. A computer system, as in claim 5, where the computer system allows one or more of the users to recover from errors during the interaction by using natural language or other modalities of interaction, based on the session context.

7. A computer system, as in claim 1, where each of the presentation managers is responsible for the presentation of all information communicated between the users and the computer system using one or more modalities of interaction, and guided by business logic governing presentation of information to the user.

8. A computer system, as in claim 7, where the modalities of interaction include one or more of the following: a keyboard input, a keyboard output, a speech input over a telephone, a speech output over a telephone, a speech input over a microphone, a speech output over a speaker, a mouse input, a pointing device input, a dataglove, and a device for translating signals into digital data.

9. A computer system, as in claim 1, where the action manager maintains communication and transactions with one or more of the following systems: an information retrieval system, a knowledge base of documents, a relational database, a directory of information, an internet site, and one or more computer systems.

10. A computer system, as in claim 1, where the taxonomical mapping process creates one or more relations between the logical forms and the content organization by associating one or more elements of the logical forms with one or more elements of the content organization of the content database.

11. A computer system, as in claim 10, where the taxonomical mapping process is guided by zero or more business logic rules.

12. A computer system, as in claim 1, where the network interface is an internet browser over the world-wide-web and the dialog manager keeps track of both the short term and long term history of a user's interactions with the computer system, and the user can ask for information or assistance using one or more modalities of interaction in an interactive dialog with the computer system.

13. A computer system, as in claim 1, where the dialog between one or more of the users and the computer system is synchronous and real time.

14. A computer system, as in claim 1, where one or more of the users can ask questions in natural language using a

18

user terminology to obtain information or execute transactions from a business.

15. A computer system, as in claim 1, where the computer system can transfer the natural language queries of one or more users to a human agent along with the history of the user interaction with the system so far and the computer system enabling the human agent to supervise the conversation and intervene if necessary.

16. A computer system, as in claim 1, where a depth of parsing of the parsing engine is adaptable, and can vary depending on one or more of the following: the taxonomical mapping, a business model, an information retrieval performance, an information retrieval confidence, and one or more actions of the presentation manager, one or more actions of the dialog manager, and one or more actions of the action manager.

17. A computer system, as in claim 16, where the parsing engine is optimized by one of the following: an automatic procedure and a manual intervention.

18. A computer system, as in claim 1, where two or more users are engaged in interaction with the computer system, and all of the users are in synchronous collaboration with each other and with the computer system by sharing the same session or sharing parts of their respective sessions.

19. computer system, as in claim 10, where the taxonomical mapping is represented as database tables.

20. A computer system, as in claim 19, where the said taxonomical mapping is represented as list of computer procedures, accessing sources of information outside of the computer system.

21. A computer system, as in claim 11, where the user can refer to pictures and content of a web page during interaction.

22. A computer system, as in claim 5, where the dialog manager supports automatic generation of said follow-up questions based one or more of the following: the session context and a domain lexicon.

23. A computer system, as in claim 22, where said follow-up questions can be generated to facilitate personalization capability.

24. A computer system, as in claim 9, where the action manager supports multimedia information retrieval by using one or more of the following to present or constrain the answers: the taxonomical mapping process, a knowledge of domain and customers, and the session context.

25. A computer system, as in claim 9, where the information retrieval is refined through the dialog based on one or more of the following: the taxonomical mapping process and session context.

26. A computer system, as in claim 9, where the communications among the presentation manager, the dialog manager and the action manager are achieved by a messaging format using one or more different transmission protocols.

27. A method executing on a computer system with one or more memories, one or more central processing units (CPU), and one or more network interfaces, the method comprising the steps of:

presenting information from the computer system to one or more users over one or more of the network interfaces and accepting queries from one or more of the users using one or more modalities over one or more of the network interfaces;

parsing one or more natural language phrases received over one or more of the network interfaces by one or more of the presentation managers, the natural language phrases being parsed into one or more logical forms, each logical form having a grammatical and structural organization;

US 6,829,603 B1

19

maintaining and directing interactive sessions between each of the users and the computer system by iteratively:

receiving logical forms from one or more of the presentation managers,

performing a discourse analysis on these logical forms, selectively generating follow-up questions, said follow-up questions being presented to a respective user after each iteration until said dialog manager determines further clarification is unneeded,

sending action requests to an action manager responsive to a determination that further clarification is unneeded,

receiving action results from the action manager, and sending presentation requests to one or more of the presentation managers;

receiving action requests from the dialog manager, performing the action, and sending the action results to the dialog manager; and

associating the logical forms of the natural language phrases to a content organization of a content database stored in one or more of the memories.

28. A computer system with one or more memories, one or more central processing units (CPU), and one or more network interfaces, the computer system comprising:

means for presenting information from the computer system to one or more users over one or more of the network interfaces and accepting queries from one or more of the users using one or more modalities over one or more of the network interfaces;

means for parsing one or more natural language phrases received over one or more of the network interfaces by one or more of the presentation managers, the natural language phrases being parsed into one or more logical forms, each logical form having a grammatical and structural organization;

means for maintaining and directing interactive sessions between each of the users and the computer system by iteratively:

receiving logical forms from one or more of the presentation managers,

performing a discourse analysis on these logical forms, selectively generating follow-up questions, said follow-up questions being presented to a respective user after each iteration until said dialog manager determines further clarification is unneeded,

sending action requests to an action manager responsive to a determination that further clarification is unneeded,

20

receiving action results from the action manager, and sending presentation requests to one or more of the presentation managers;

means for receiving action requests from the dialog manager, performing the action, and sending the action results to the dialog manager; and

means for associating the logical forms of the natural language phrases to a content organization of a content database stored in one or more of the memories.

29. A computer program product for use on a computer system which causes the computer system to perform the steps of:

presenting information from the computer system to one or more users over one or more of the network interfaces and accepting queries from one or more of the users using one or more modalities over one or more of the network interfaces;

parsing one or more natural language phrases received over one or more of the network interfaces by one or more of the presentation managers, the natural language phrases being parsed into one or more logical forms, each logical form having a grammatical and structural organization;

maintaining and directing interactive sessions between each of the users and the computer system by iteratively:

receiving logical forms from one or more of the presentation managers,

performing a discourse analysis on these logical forms, selectively generating follow-up questions, said follow-up questions being

presented to a respective user after each iteration until said dialog manager

determines further clarification is unneeded.

sending action requests to an action manager responsive to a determination that further clarification is unneeded,

receiving action results from the action manager, and sending presentation requests to one or more of the presentation managers;

receiving action requests from the dialog manager, performing the action, and sending the action results to the dialog manager; and

associating the logical forms of the natural language phrases to a content organization of a content database stored in one or more of the memories.

* * * * *

EXHIBIT G

WEBSTER'S NEW WORLDTM

Telecom Dictionary

Ray Horak



Wiley Publishing, Inc.

Webster's New World® Telecom Dictionary

Published by

Wiley Publishing, Inc.

10475 Crosspoint Boulevard

Indianapolis, IN 46256

www.wiley.com

Copyright © 2008 by Wiley Publishing, Inc., Indianapolis, Indiana

Published simultaneously in Canada

ISBN: 978-0-471-77457-0

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Library of Congress Cataloging-in-Publication Data

Horak, Ray.

Webster's New World telecom dictionary / Ray Horak.

p. cm.

ISBN 978-0-471-77457-0 (pbk.)

1. Telecommunication—Dictionaries. I. Title.

TK5102.H65 2007

621.38203—dc22

2007024232

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600. Requests to the Publisher for permission should be addressed to the Legal Department, Wiley Publishing, Inc., 10475 Crosspoint Blvd., Indianapolis, IN 46256, (317) 572-3447, fax (317) 572-4355, or online at <http://www.wiley.com/go/permissions>.

Limit of Liability/Disclaimer of Warranty: The publisher and the author make no representations or warranties with respect to the accuracy or completeness of the contents of this work and specifically disclaim all warranties, including without limitation warranties of fitness for a particular purpose. No warranty may be created or extended by sales or promotional materials. The advice and strategies contained herein may not be suitable for every situation. This work is sold with the understanding that the publisher is not engaged in rendering legal, accounting, or other professional services. If professional assistance is required, the services of a competent professional person should be sought. Neither the publisher nor the author shall be liable for damages arising herefrom. The fact that an organization or Website is referred to in this work as a citation and/or a potential source of further information does not mean that the author or the publisher endorses the information the organization or Website may provide or recommendations it may make. Further, readers should be aware that Internet Websites listed in this work may have changed or disappeared between when this work was written and when it is read.

For general information on our other products and services please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Trademarks: Wiley, the Wiley logo, Webster's New World, the Webster's New World logo, We Define Your World, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates, in the United States and other countries, and may not be used without written permission. All other trademarks are the property of their respective owners. Wiley Publishing, Inc., is not associated with any product or vendor mentioned in this book.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

P.861 The ITU-T standard defining the Perceptual Speech Quality Measurement (PSQM) method of objectively evaluating voice transmission quality. PSQM is an automated method developed to measure the perceived quality of real-time voice as impacted by compression codecs. An enhanced version known as PSQM+ was later developed to measure the effects of packet loss and other network impairments. PSQM compares a distorted voice sample to an original, clear voice sample and uses a complex analytical process to evaluate the difference in terms of factors that influence the perceptions of human listeners. The ultimate distortion score corresponds closely with the mean opinion score (MOS) yielded by the panel of human listeners employed using the P.800 method. P.861 was withdrawn and replaced by P.862, defining the Perceptual Evaluation of Speech Quality (PESQ). See also *codec*, *compression*, *MOS*, *P.800*, *P.862*, and *PSQM*.

P.862 The ITU-T standard defining the Perceptual Evaluation of Speech Quality (PESQ) method of objectively evaluating transmission quality. A companion to the more subjective P.800, PESQ is an automated method that addresses the effects of filters, jitter, and coding distortions. PESQ replaced the Perceptual Speech Quality Measurement (PSQM) method, which was viewed as limited in certain applications. PESQ is considered to be an intrusive test method as it inserts a reference signal into the device under test. See also *ITU-T*, *P.800*, *PSQM*, and *toll quality*.

PABX (Private Automatic Branch eXchange) Generally synonymous in contemporary usage with private branch exchange (PBX), PABX refers to an automatic PBX, as compared to a manual PBX, or cordboard. The term *PBX* is preferred in North America, and PABX in much of the balance of the world. See *PBX* for more detail.

packet 1. In the generic sense, referring to the manner in which data are organized into discrete units for transmission and switching through a data network. The data unit can be known as a block, frame, cell, or packet, depending on the protocol specifics. The packet comprises a header, payload, and sometimes a trailer, again depending on protocol specifics. The packet can be a user packet containing user data, or a signaling and control packet for various network monitoring, alerting and alarming, maintenance, and other administrative purposes. The payload can be a complete message, a fragment or segment of a message, or an aggregation of bits or bytes that form a short portion of a long data stream associated with a voice or video call. See also *bit*, *block*, *byte*, *cell*, *data stream*, *fragment*, *frame*, *header*, *message*, *payload*, *protocol*, *segment*, and *trailer*. **2.** In a technology-specific sense, a packet is a data unit in an internetwork, such as the Internet or other packet-switched network in which routers interconnect networks and subnetworks to exchange traffic between nodes. In terms of the OSI Reference Model, a packet is defined in Layer 3, the Network Layer. Blocks, cells, and frames are defined in Layer 2, the Data Link Layer, and have local significance, only. See also *block*, *cell*, *datagram*, *Data Link Layer*, *frame*, *Internet*, *Network Layer*, *OSI Reference Model*, *packet switch*, and *router*.

packet assembler/disassembler (PAD) See *PAD*.

packet-filtering firewall A security firewall that examines all data packets, forwarding or dropping individual packets based on predefined rules that specify where a packet is permitted to go, in consideration of both the authenticated identification of the user and the originating address of the request. See also *authentication*, *firewall*, *proxy firewall*, *security*, and *stateful inspection firewall*.

packet Internet groper (ping) See *ping*.

packet layer protocol (PLP) See *PLP*.

Packet over SONET (POS) See *POS*.

packet switch A device that switches data organized into packets, discrete sets of data that may take the specific form of packets, frames, or cells depending on the network technology specifics. For example, packet switches switch packets in networks based on the Internet Protocol (IP), frames in networks based on the Frame Relay or Ethernet protocol, and cells in those based on the Asynchronous Transfer Mode



US006081907A

United States Patent [19]
Witty et al.

[11] **Patent Number:** **6,081,907**
[45] **Date of Patent:** **Jun. 27, 2000**

[54] **DATA DELIVERY SYSTEM AND METHOD FOR DELIVERING DATA AND REDUNDANT INFORMATION OVER A UNIDIRECTIONAL NETWORK**

[75] Inventors: **Carl R. Witty; Kenneth J. Birdwell; James Randall Sargent**, all of Bellevue; **Brian Moran**, Issaquah, all of Wash.

[73] Assignee: **Microsoft Corporation**, Redmond, Wash.

[21] Appl. No.: **08/871,657**

[22] Filed: **Jun. 9, 1997**

[51] **Int. Cl.⁷** **G06F 11/00**

[52] **U.S. Cl.** **714/6; 714/752**

[58] **Field of Search** 395/182.04, 182.09, 395/182.05, 182.14; 370/216, 218, 389, 409; 709/200, 201, 203, 212, 213, 227, 232, 238; 714/6, 18, 20, 716, 717, 761, 762, 755, 752

[56] **References Cited**

U.S. PATENT DOCUMENTS

4,769,818	9/1988	Mortimer	714/755
5,335,277	8/1994	Harvey et al.	380/20
5,359,367	10/1994	Stockhill	348/552
5,559,808	9/1996	Kostreski et al.	370/108
5,565,909	10/1996	Thibadeau et al.	348/9
5,625,864	4/1997	Budow et al.	455/4.2
5,627,829	5/1997	Gleeson et al.	370/230
5,650,831	7/1997	Farwell	348/734
5,666,293	9/1997	Metz et al.	395/200.5
5,684,807	11/1997	Bianchini et al.	395/182.02
5,706,048	1/1998	Davis	348/12

5,727,065	3/1998	Dillon	380/49
5,727,865	3/1998	Dillon	380/49
5,778,187	7/1998	Monteiro et al.	395/200.61
5,809,543	9/1998	Byers et al.	395/182.04
5,854,897	12/1998	Radziewicz et al.	395/200.54
5,859,852	1/1999	Moura et al.	370/449

OTHER PUBLICATIONS

McCarthy, Michael, "VRML 97 Mbone Broadcast Announcement," Feb. 1997.

The Mbone Information Web @ www.mbone.com.

Kashpureff.ORG Network Information Center, "Search Internet Drafts," @ www.kashpureff.org.

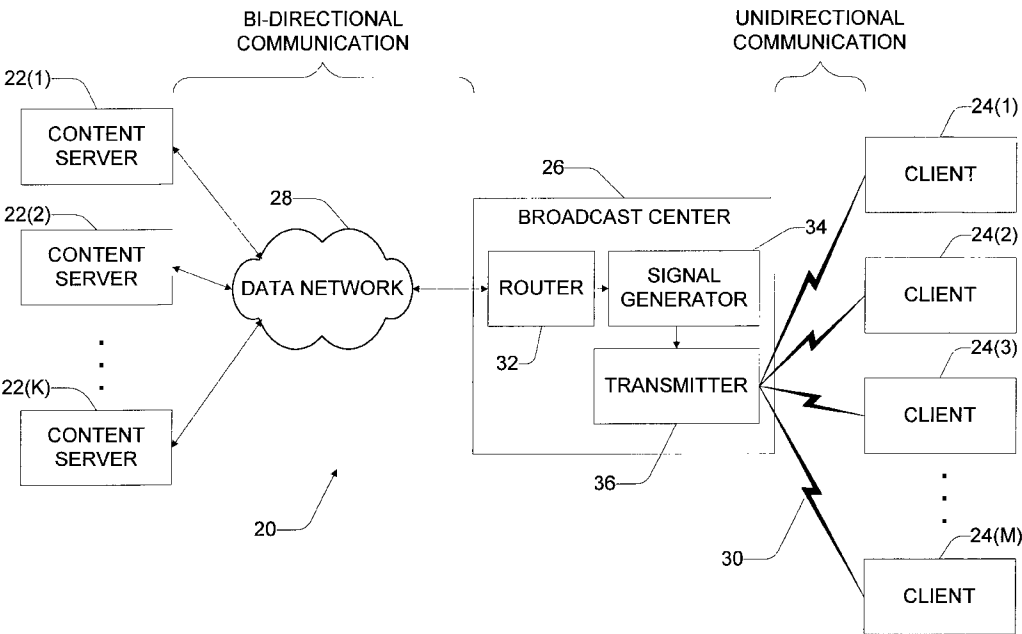
R. Panabaker et al. "The Transmission of IP Over the Vertical Blanking Interval Of a Television Signal"—Internet Reference, [HTTP://www.microsoft.com](http://www.microsoft.com), 1997.

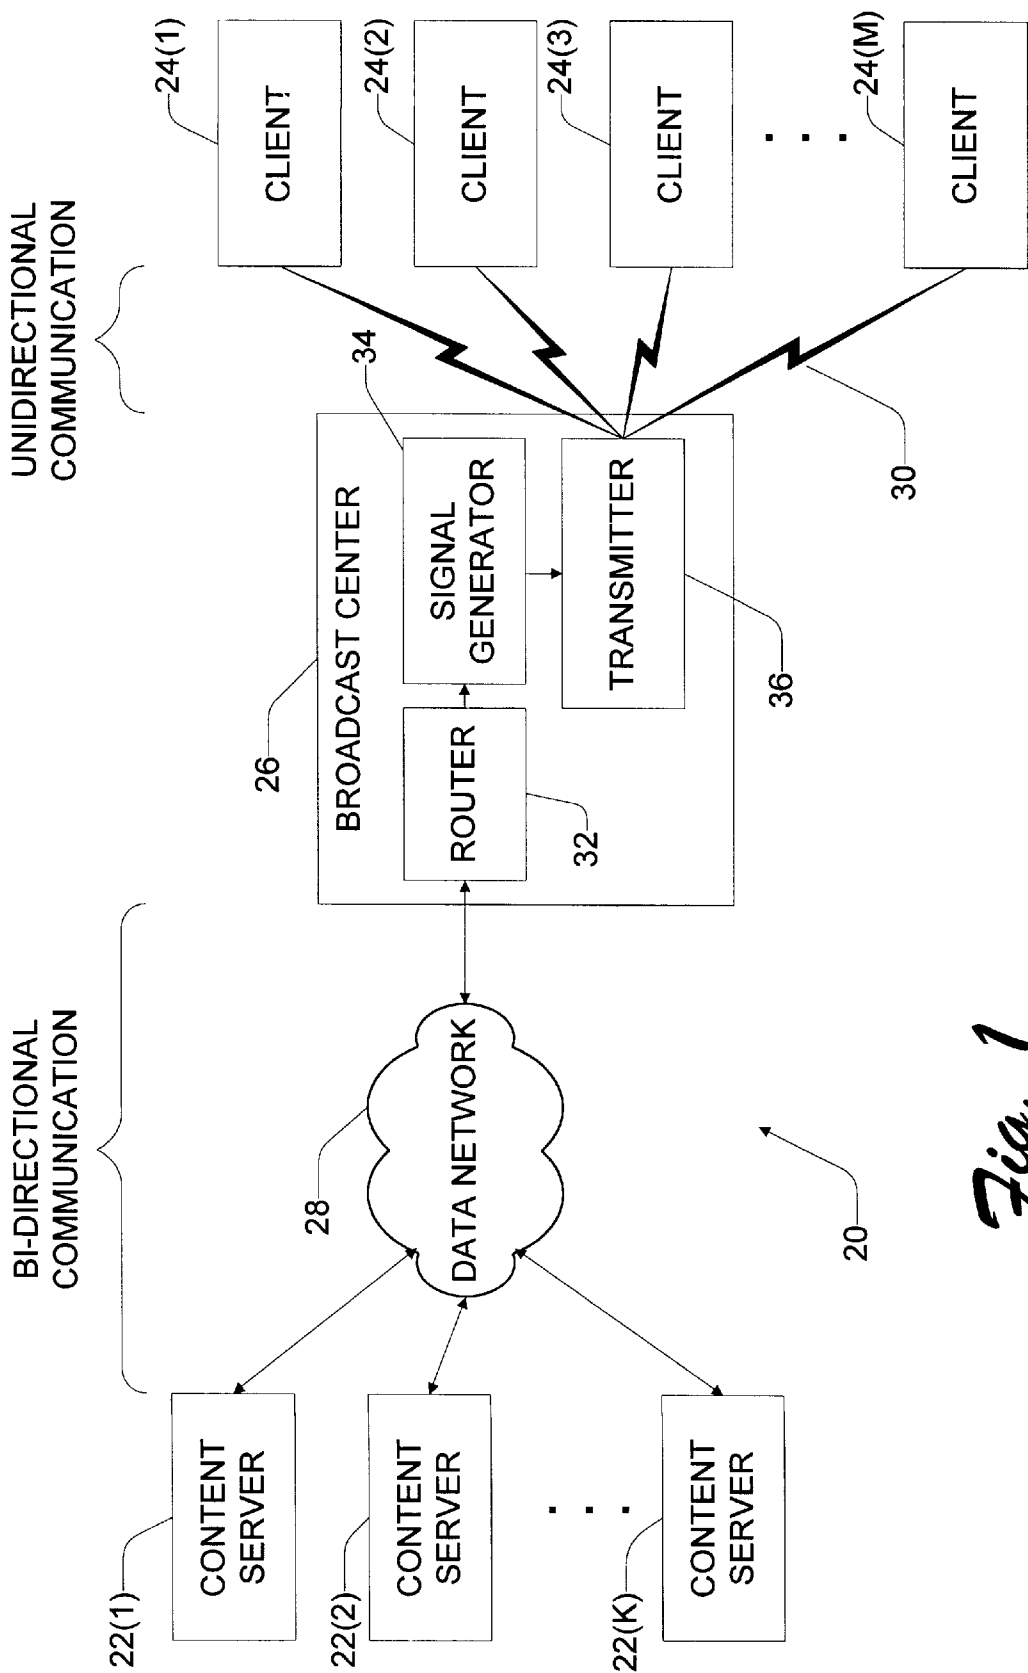
Primary Examiner—Robert W. Beausoliel, Jr.
Assistant Examiner—Pierre Eddy Elisea
Attorney, Agent, or Firm—Lee & Hayes, PLLC

[57] **ABSTRACT**

A data delivery system facilitates transmission of data packets from a content server to multiple clients over a unidirectional network. A redundancy formatter resident at the server groups multiple data packets into a redundancy group and generates at least one redundancy packet containing redundancy information derived from the data packets in the redundancy group. The data packets and redundancy packet are sent over the unidirectional network to the client. In the event that a packet is lost, a packet rebuilder resident at each client reconstructs the missing data packet from the successfully transmitted data packets in the redundancy group and the redundancy packet for the redundancy group.

35 Claims, 7 Drawing Sheets








Fig. 2

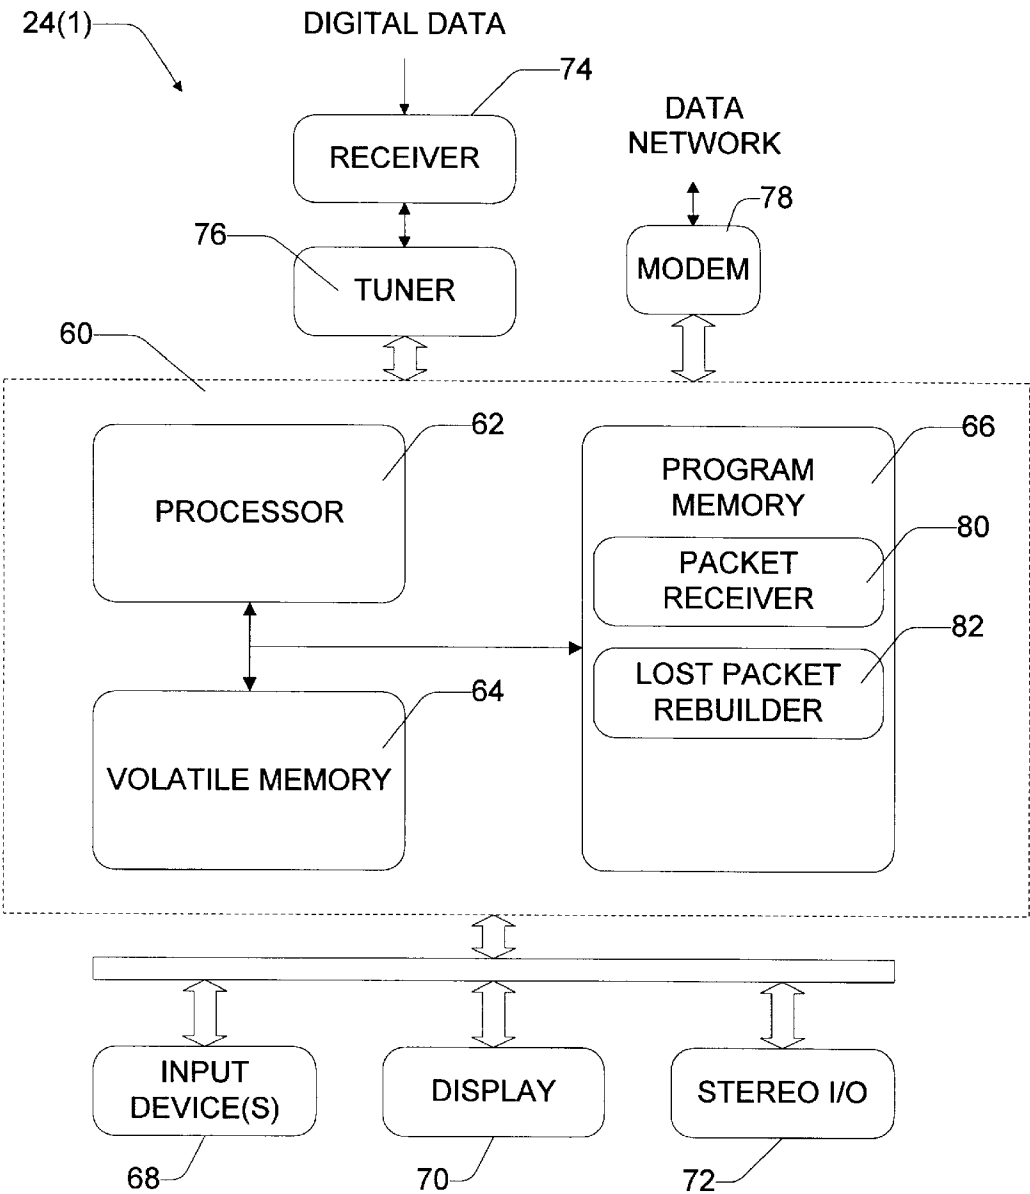


Fig. 3




Fig. 4

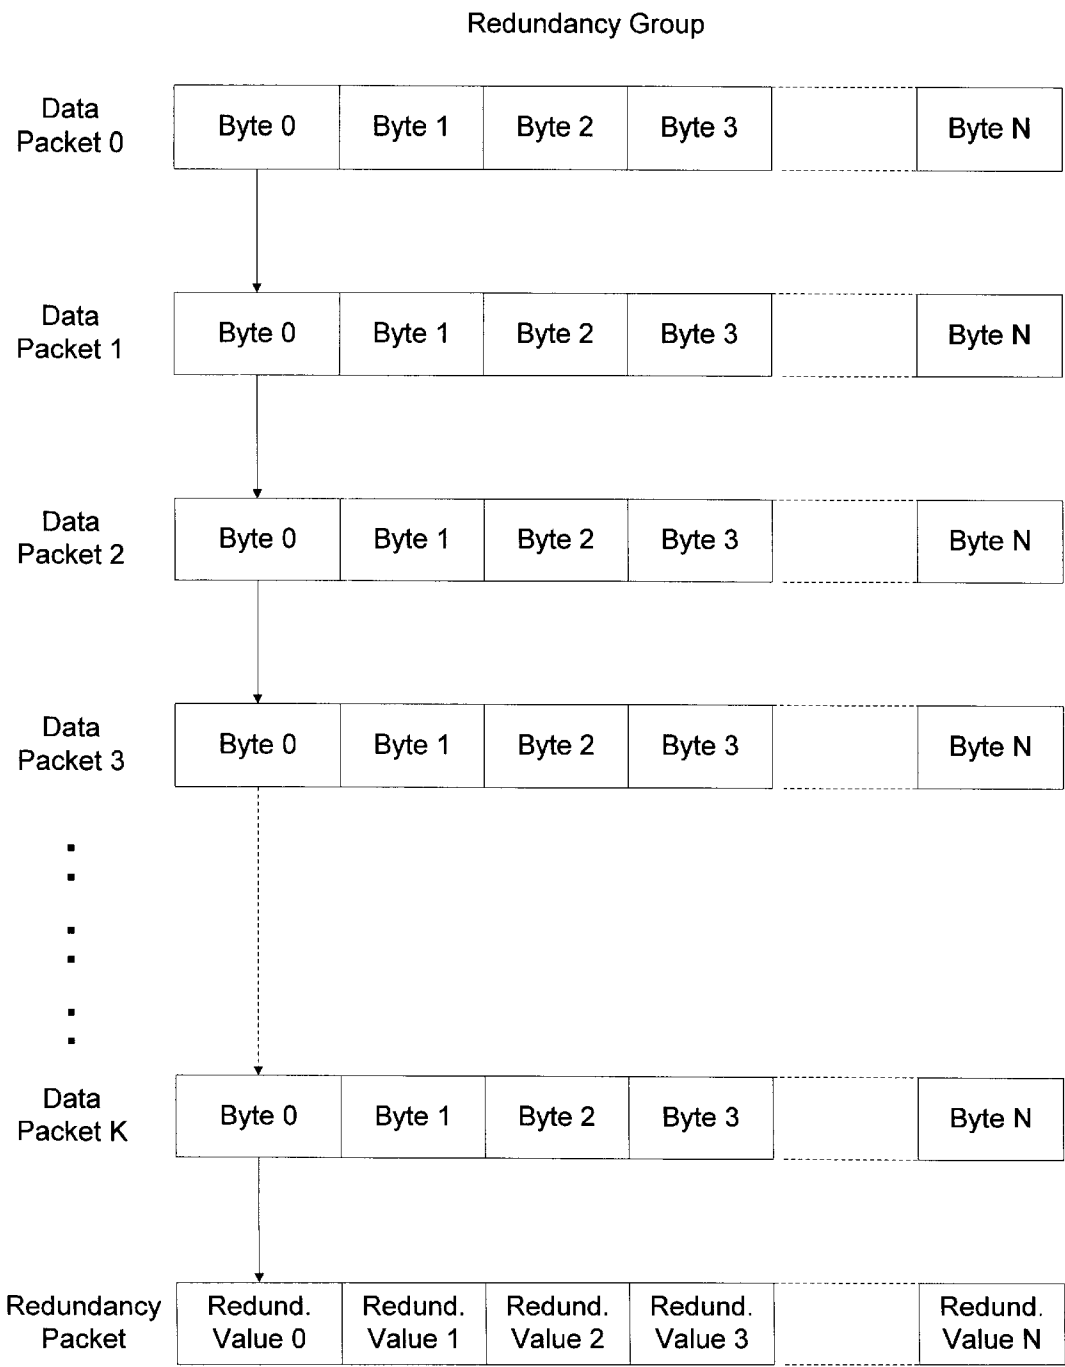



Fig. 5

Fig. 6






Fig. 7

6,081,907

1

**DATA DELIVERY SYSTEM AND METHOD
FOR DELIVERING DATA AND REDUNDANT
INFORMATION OVER A UNIDIRECTIONAL
NETWORK**

TECHNICAL FIELD

This invention relates to a data delivery system in which computer data and other content are served from one or more servers over a unidirectional network to one or more clients. More particularly, this invention relates to a data delivery system and method for delivering redundant information along with the data over the unidirectional network.

BACKGROUND OF THE INVENTION

Conventional computer networks are bi-directional, allowing data communication in both directions between servers and clients. Transmitting data over these bi-directional data networks has been a mainstay of computer technology for many years and the communication protocols are well established. Under conventional communication protocols, it is common for the client to initiate connection with the server and to request desired data from the server. As part of the request, the client sends information pertaining to how the data should be sent. For example, the client might include a client address, TCP port number, and so forth.

Digital data, whether transmitted over a wire-based distribution network (e.g., local area network, wide area network, cable, etc.) or a wireless distribution network (e.g., satellite, RF, paging, etc.), is typically packetized and sent over the network in individual packets. As the data packets are sent, both the server and client carefully track every packet. If a packet is lost during transmission, the server and client contact one another and the lost packet is resent.

Apart from the classic bi-directional data networks, there is an increasing interest in the use of broadcast or multicast networks to deliver computer data and other content to clients. These types of distribution networks are unidirectional in that data flows from the server to the clients, but no return communication is possible over the same communication path. As a result, the broadcast network cannot support the common protocols used for two-way communication over a bi-directional network, such as client-driven connections and data requests, because the clients are unable to communicate over the broadcast communication link to the server.

One problem that arises in the transmission of data packets over a unidirectional network is the occasional loss of data packets during transmission. The absence of a return communication path from the client to the server over the unidirectional network renders it impossible for the client to ask the server to resend the missing packet. As a result, the server blindly sends the packets over the unidirectional network and hopes the client receives all of the sent packets. The problem is amplified in broadcast or multicast environments where the server is sending the packets out to many clients, all of whom may be dropping different data packets.

The inventors have developed a system and method which address this problem.

SUMMARY OF THE INVENTION

A data delivery system facilitates transmission of data packets from a content server to multiple clients over a unidirectional network. A redundancy formatter resident at the server groups multiple data packets into a redundancy

2

group and generates at least one redundancy packet containing redundancy information derived from the data packets in the redundancy group. The data packets and redundancy packet are sent over the unidirectional network to the client. In the event that a packet is lost, a packet rebuilder resident at each client reconstructs the missing data packet from the successfully transmitted data packets in the redundancy group and the redundancy packet for the redundancy group.

According to an aspect of this invention, each packet has a header comprising a unique packet identifier, a content identifier to uniquely identify the data packet as belonging to a particular data transmission, a redundancy group identifier to identify the data packet as belonging to a particular redundancy group, and a redundancy group count of the total number of packets contained within the redundancy group. To protect against loss of multiple packets in a sequence, packets from different redundancy groups are interleaved during transmission over the unidirectional network.

Upon receiving the data and redundancy packets, the client examines the packet headers to gather packets with the same redundancy group identifier. For each redundancy group, the client detects whether any data packet is missing from the redundancy group by comparing the redundancy group count in a packet header with the actual number of gathered packets. If a packet is missing, the client regenerates the missing packet from the other packets in the redundancy group.

According to another aspect of this invention, each data packet is configured with multiple bytes at predefined positions. The redundancy packet is created with multiple bytes which are computationally derived (e.g., using an XOR function) from corresponding bytes in the same byte positions of the data packet bytes. The same computation of corresponding bytes is used to reconstruct a missing packet.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a diagrammatic illustration of a data delivery system for sending data packets and redundancy information over a unidirectional network.

FIG. 2 is a block diagram of a server computing unit.

FIG. 3 is a block diagram of a client computing unit.

FIG. 4 is a flow diagram showing steps in a method for serving data packets over a unidirectional network.

FIG. 5 is a diagrammatic illustration of multiple data packets within a redundancy group. FIG. 5 shows one example of how a redundancy packet is generated from the multiple data packets.

FIG. 6 is a diagrammatic illustration of a data packet structure.

FIG. 7 is a flow diagram showing steps in a method for receiving data packets transmitted over a unidirectional network.

**DETAILED DESCRIPTION OF THE
PREFERRED EMBODIMENT**

This invention concerns delivery of computer data and other content over a unidirectional network. One or more content servers are connected to serve the data over the unidirectional network to one or more clients. The unidirectional network may be the only distribution means between the server and client, or it may be a single link in a bigger distribution network. Generally, the unidirectional network is a network or link that prohibits a client from contacting the server via the same communication path over which the data is being served.

FIG. 1 shows an exemplary data delivery system 20 in which computer data and other content is delivered from multiple content servers 22(1), 22(2), . . . , 22(K) to multiple clients 24(1), 24(2), 24(3), . . . , 24(M). In this implementation, the content servers 22(1)–22(K) are connected to a broadcast center 26 via a bi-directional data network 28 which enables two-way communication between the content servers 22(1)–22(K) and the broadcast center 26. The content servers serve data in the form of audio, video, animation, bit maps or other graphics, applications or other executable code, text, hypermedia, or other multimedia types.

The bi-directional data network 28 represents various types of networks, including the Internet, a LAN (local area network), a WAN (wide area network), and the like. The data network 28 can be implemented in a number of ways, including wire-based technologies (e.g., fiber optic, cable, wire, etc.) and wireless technologies configured for two-way communication (e.g., satellite, RF, etc.). The data network 28 can further be implemented using various available switching technologies (e.g., ATM (Asynchronous Transfer Mode), Ethernet, etc.) and different data communication protocols (e.g., IP, IPX, etc.).

The broadcast center 26 receives the data served from the content servers 22(1)–22(K) over the network 28 and broadcasts the data over a broadcast network 30 to the clients 24(1)–24(M). In the FIG. 1 implementation, the unidirectional network is embodied as the broadcast network 30 in which the data is carried in one direction from the broadcast center 26 to the many clients 24(1)–24(M). The clients are unable to reply or initiate communication to the broadcast center 26 using the broadcast network 30.

The broadcast network 30 can be implemented in a variety of ways. For instance, the broadcast network might be implemented as a wireless network configured for one-way transmission (i.e., satellite, radio, microwave, etc.). The broadcast network might also be a network which supports two-way communication, but is predominately used for unidirectional multicasting from the broadcast center 26 to many clients simultaneously. Although only one broadcast center 26 is illustrated for explanation purposes, the system 20 can scale to include multiple broadcast centers coupled between numerous servers 22 and numerous clients 24.

The broadcast center 26 includes a router 32, a signal generator 34, and a broadcast transmitter 36. The router 32 is coupled to the bi-directional data network 28 to receive the data served over the network 28 from the content servers 22(1)–22(K). The router 32 is a final node of the data network 28 in which data communication is bi-directional to that point and unidirectional past that point. The router 32 is preferably configured as a bridge-router between the traditional data network 28 and the broadcast network 30. A bridge-router is capable of supporting video and audio broadcast transmission.

Data is received at the router 32 and converted from the network packet format to a format appropriate for broadcast transmission. The signal generator 34 generates a broadcast signal with the data embedded thereon to carry the data over the broadcast network 30. The broadcast signal is passed to the transmitter 36 where it is broadcast over the broadcast network 30 to the clients 24(1)–24(M). The clients might still be able to communicate with the broadcast center 26 or content servers 22(1)–22(K) using a different back channel, but this aspect is not shown in the drawings.

FIG. 2 shows an exemplary implementation of a content server 22(1). It includes a server computer 40 having a

processor 42 (e.g., Pentium® Pro microprocessor from Intel Corporation), volatile memory 44 (e.g., RAM), and program memory 46 (e.g., ROM, disk drive, floppy disk drive, CD-ROM, etc.). The computer 40 is configured, for example, as a personal computer or workstation running a multitasking, disk-based operating system, such as Windows® NT from Microsoft Corporation. The server computer 40 is connected to the data network 28 via a network connection 48. The content server 22(1) has multiple storage disks 50, which might be implemented as a disk array, to store various forms of content. In this illustration, the content server 22(1) is shown configured as continuous media file server which serves data files from a disk array of storage disks 50.

The server 22(1) is illustrated with two software programs: a packet encoder 52 and a redundancy formatter 54. Each program is stored in program memory 46, loaded into volatile memory 44 when launched, and executed on the processor 42. The redundancy formatter 54 groups multiple data packets into a redundancy group and generates at least one redundancy packet containing redundancy information derived from the data packets in the redundancy group. The packet encoder 52 encapsulates packets of data with appropriate headers for transmission over the data network and broadcast network. The operation of the redundancy formatter 54 and packet encoder 52 is described below in more detail with reference to FIGS. 4–6.

With reference again to FIG. 1, the clients 24(1)–24(M) can be implemented in a number of ways, including desktop computers, laptop computers, set-top boxes, and computer enhanced television units. As an example implementation, the client is a broadcast-enabled personal computer.

FIG. 3 shows an exemplary configuration of a client 24(1) implemented as a broadcast-enabled computer. It includes a central processing unit 60 having a processor 62 (e.g., x86 or Pentium® microprocessor from Intel Corporation), volatile memory 64 (e.g., RAM), and program memory 66 (e.g., ROM, disk drive, floppy disk drive, CD-ROM, etc.). The client 24(1) has one or more input devices 68 (e.g., keyboard, mouse, etc.), a computer display 70 (e.g., VGA, SVGA), and a stereo I/O 72 for interfacing with a stereo audio system.

The client 24(1) includes a digital broadcast receiver 74 (e.g., satellite dish receiver, RF receiver, microwave receiver, NTSC-VBI receiver, digital cable TV decoder, etc.) and a tuner 76 which tunes to appropriate frequencies or addresses of the broadcast network 30. The tuner 76 is configured to receive digital broadcast data in a particularized format, such as MPEG-encoded digital video and audio data, as well as digital data in many different forms, including software programs and programming information in the form of data files. The client 24(1) also has a modem 78 which provides dial-up access to the data network to provide a back channel or secondary link to the content servers. In other implementations of a secondary link, the modem 78 might be replaced by a network card, or an RF receiver, or other type of port/receiver which provides access to the secondary link.

The client 24(1) runs an operating system which supports multiple applications. The operating system is preferably a multitasking operating system which allows simultaneous execution of multiple applications. The operating system employs a graphical user interface windowing environment which presents the applications or documents in specially delineated areas of the display screen called “windows.” One preferred operating system is a Windows® brand oper-

ating system sold by Microsoft Corporation, such as Windows® 95 or Windows® NT or other derivative versions of Windows®. It is noted, however, that other operating systems which provide windowing environments may be employed, such as the Macintosh operating system from Apple Computer, Inc. and the OS/2 operating system from IBM.

One example implementation of a broadcast-enabled PC is described in a co-pending U.S. patent application Ser. No. 08/503,055, entitled "Broadcast-Enabled Personal Computer," filed Jan. 29, 1996 in the names of Gabe L. Newell, Dan Newell, Steven J. Fluegel, David S. Byrne, Whitney McCleary, James O. Roberts, Brian K. Moran; William B. McCormick, T. K. Backman, Kenneth J. Birdwell, Joseph S. Robinson, Alonzo Garipey, Marc W. Whitman, and Larry Brader. This application is assigned to Microsoft Corporation, and is incorporated herein by reference.

The client 24(1) is illustrated with two software programs: a packet receiver 80 and a lost packet rebuild 82. Each program is stored in program memory 66, loaded into volatile memory 64 when launched, and executed on the processor 62. The packet receiver 80 decodes the data packets received by the digital receiver 74. The packet receiver 80 performs various error correction procedures on the packets, obtains information from the headers, and extracts the data payloads for local use. The lost packet rebuild 82 operates in conjunction with the packet receiver 80 to reconstruct any missing data packet using the successfully transmitted data packets in the missing data packet's redundancy group and the redundancy packet for the redundancy group. The operation of the packet receiver 80 and the lost packet rebuild 82 is described below in more detail with reference to FIG. 7.

FIG. 4 shows exemplary steps in a method for serving data packets over the unidirectional network. Software (e.g., packet encoder 52 and redundancy formatter 54) executing at the content servers performs these steps. At step 90, the content server forms individual packets of data. These packets have a finite number of bytes. Preferably, the packets are configured to comply with a multicast data protocol, such as multicast UDP (User Datagram Protocol) which is a standard network protocol.

At step 92 in FIG. 4, the packet encoder 52 assigns a content identifier to each data packet to identify the data packet as belonging to a particular data transmission. The data delivery system is designed to support multiple simultaneous transmissions from different content servers. The content identifiers in the data packets distinguish the data packets from among the different broadcast transmissions. For short data transmissions, the content identifier is a unique value derived from the actual content to be transmitted. A hashing algorithm (e.g., MD5, SHA, etc.) computes a digest of the actual data which is used as the content ID. For longer data transmissions in which it is infeasible to compute a hash of the actual data, the content identifier is a large unique random number that distinguishes packets in one data transmission from packets in other data transmissions.

At step 94 in FIG. 4, the packet encoder 52 assigns a unique packet identifier to each data packet to differentiate individual packets within the same data transmission from one another.

At step 96, the redundancy formatter 54 groups data packets into multiple redundancy groups (e.g., M redundancy groups). The number of redundancy groups varies

depending upon the size of the transmission and the anticipated percentage of lost packets. Each redundancy group has multiple data packets. The number of packets in each group also varies according to a number of factors, such as size of transmission, anticipated number of lost packets, redundancy computing efficiency, transmission bandwidth, and so forth.

The redundancy formatter 54 processes each redundancy group (step 98 in FIG. 4) to generate at least one redundancy packet containing redundancy information for that redundancy group. As one example implementation, the redundancy formatter derives a redundancy packet by computing the XOR function of the same byte position from each data packet in the redundancy group (step 100 in FIG. 4).

FIG. 5 shows the byte-wise technique for generating a redundancy packet from multiple data packets within a redundancy group. The illustrated redundancy group has data packets 0, 1, 2, 3, . . . , K. Each data packet has N bytes. The bytes of each data packet correspond to one another by position, with all bytes 0 being present in the first position of their respective data packets, all bytes 1 being present in the second position, and so on. The redundancy formatter 54 computes a redundancy value for each corresponding byte position of the redundancy data packet. The first byte position in the redundancy packet—redundant value 0—is derived by XORing all of the bytes 0 from data packets 0–K. The second byte position in the redundancy packet—redundant value 1—is derived by XORing all of the bytes 1 from data packets 0–K, and so on.

FIG. 5 is illustrated for example purposes. Rather than a byte-wise approach, the redundancy values may be computed using multi-byte words, or an arbitrary segment size. Other computations may be used in place of the XOR function, including Hamming codes, Forward Error Correction codes, etc. Additionally, depending on the error correction or erasure correction algorithm chosen, more than one redundancy packet may be formed for each redundancy group.

With continuing reference to step 102 to FIG. 4, the redundancy formatter 54 assigns a redundancy group identifier to all K data packets within the redundancy group to distinguish those data packets from packets belonging to other redundancy groups. The redundancy packet for the redundancy group is also tagged with the redundancy group identifier. The redundancy group identifier can be a number, alpha string, or other type of indicia that uniquely differentiates packets in different groups. At step 104, the redundancy formatter 54 assigns a redundancy group count to all K data packets and the redundancy packet within the redundancy group. The redundancy group count is the total number of packets (including data and redundancy packets) within the redundancy group, which is this example is K+1. This redundancy group count is used by the receiver to determine if any packets are missing from the redundancy group.

At step 106 in FIG. 4, the redundancy formatter 54 proceeds to the next redundancy group to generate a redundancy packet based on the data packets within that redundancy group and to mark all packets with the appropriate identifiers.

As each redundancy group is completed, or after the entire data transmission is encoded, the packets are served over the unidirectional network to the clients. Both the data packets and the redundancy packets are sent. Typically, any lost data packets are lost in bursts during the broadcast transmission. To minimize loss of multiple packets in a sequence, data

packets from different redundancy groups are interleaved. The data packets are often interleaved naturally as a product of simultaneous multicast of entirely different content served over the unidirectional network from different servers. The content identifiers and redundancy group identifiers are used by the clients to distinguish the packets, as is described below in more detail.

FIG. 6 shows an exemplary data structure 110 for a data packet formed by the packet encoder 52 and redundancy formatter 54. The data structure 110 has a header 112 and a data payload 114 concatenated together. The data payload 112 holds J bytes of actual data content, where the number J depends upon the transmission protocol.

The header 112 includes two MFTP (Multi File Transport Protocol) headers 116 and 118, an error correction protocol header 120, and a group header 122. The first MFTP header 116 is attached to the data packet 114. The MFTP header 116 contains a four-byte field “high_len” which stores the high order 32 bits of total data length; a four-byte field “low_len” which holds the low order 32 bits of total data length; a four-byte field “high_pos” which holds the high order 32 bits of current data position; a four-byte field “low_pos” which holds the low order 32 bits of current data position; and a one-byte field “data_size” indicating the size of the attached data payload 114.

The second MFTP header 118 is wrapped around the data payload 114 and first MFTP header 116. The second MFTP header 118 contains a one-byte field “version” which identifies the current version of the MFTP protocol; a one-byte field “flags1” of undefined flags; a two-byte field “flags2” of undefined flags; and a one-byte field “mftp_data” which holds data for specific MFTP protocol.

Following the two MFTP headers, the packet is further wrapped with an ECC header 120 provided by the redundancy formatter 54. The ECC header contains information pertaining to the XOR ECC protocol used to develop the redundancy packet. The ECC header 120 contains a four-byte redundancy group ID field to identify the packet as belonging to a particular redundancy group. The ECC header 120 further includes a one-byte field which holds the redundancy group count. This count specifies the total number of packets contained within the redundancy group identified by the group ID field. A one-byte packet ID field contains the packet identifier which is unique for each packet. The remaining fields of the ECC header 120 are a one-byte field to hold the number of different redundancy groups that are interleaved; a one-byte field of undefined flags; and a one-byte field for MFTP protocol information.

The first three headers 116–120 are unique to each packet, as the values within the headers change from packet to packet.

The group header 122, which wraps the data payload 114 and previous headers 116–120, is common for all packets within the same transmission. The group header 122 contains a four-byte CRC field which contains CRC (cyclical redundancy check) information for the entire packet. The group header 122 contains a two-byte ECP field which specifies the error correction protocol used to derive the redundancy packet and form the ECC header 120. A one-byte “flags” field is presently unused, but may be defined later. A one-byte “cont.len” field specifies the length of the content identifier, which is held in the neighboring L-byte “cont.ID” field. As described above, the content identifier stored in the cont.ID field is a hash digest or other unique representation of the actual data content of the entire data transmission. The last field is a four-byte field which contains protocol-related information.

FIG. 7 shows exemplary steps in a method for receiving data packets transmitted over a unidirectional network. Software (e.g., packet receiver 80 and lost packet rebuilder 82) executing at the clients performs these steps. At step 130, the client receives the packets from the unidirectional network. The client is configured to handle the multicast data protocol, such as multicast UDP. At step 132, the packet receiver 80 examines the group header 122 and matches data packets having the same content identifier. This process separates the data packets into their respective transmissions.

At step 134 in FIG. 7, the packet receiver 80 examines the ECC header 120 and gathers packets with the same redundancy group identifier. This process differentiates packets within the same transmission according to their redundancy group.

Once the packets are gathered into redundancy groups, the lost packet rebuilder 82 examines each redundancy group (step 136 in FIG. 7). For each redundancy group, the lost packet rebuilder 82 examines the ECC headers 120 of the packets to detect if any packet is missing (step 138 in FIG. 7). The ECC header 120 contains the redundancy group count specifying the number of packets contained within the redundancy group. The lost packet rebuilder 82 detects whether any data packet is missing from a particular redundancy group by comparing the redundancy group count in the ECC headers with the actual number of packets received from the unidirectional network. If no packet is missing (i.e., the “no” branch from step 138), the lost packet rebuilder 82 continues to the next redundancy group (step 142).

Conversely, if a packet is missing (i.e., the “yes” branch from step 138), the lost packet rebuilder 82 reconstructs the missing data packet using the successfully received data packets and the redundancy packet of the corresponding redundancy group (step 140 in FIG. 7). More particularly, the same XOR function is computed for each byte position of the available packets to reconstruct the bytes of the missing packet. The lost packet rebuilder 82 might optionally attempt to determine which data packet is missing based on the unique packet identifiers of the received packets. After the data packet is recovered, the lost packet rebuilder proceeds to the next redundancy group (step 142).

When using the XOR function, only one packet from each redundancy group can be lost and regenerated. When two or more packets are lost from the same redundancy group, the packets cannot be recovered using the XOR function. Accordingly, if a high percentage of packets are expected to be lost during a transmission, the server may decide to form a greater number of redundancy groups, with each group having a smaller number of packets. In this manner, there is much less chance of losing multiple packets per group.

With other functions, however, more than one packet may be lost without loss of data content. For example, using Hamming codes and related FEC algorithms, multiple lost packets can be reconstructed depending on the amount of redundant information sent (assuming uncompressible data content the ratio of redundant data required to lost packet recovery will be at least 1 to 1).

In compliance with the patent statute, the invention has been described in language more or less specific as to structural and methodical features. It is to be understood, however, that the invention is not limited to the specific features described, since the means herein disclosed comprise preferred forms of putting the invention into effect. The invention is, therefore, claimed in any of its forms or modifications within the proper scope of the appended

6,081,907

9

claims appropriately interpreted in accordance with the doctrine of equivalents.

What is claimed is:

1. A method for sending data packets over a unidirectional network, comprising the following steps:

grouping multiple data packets into a redundancy group;
forming at least one redundancy packet containing redundancy information generated from the data packets in the redundancy group; and

sending the data packets and the redundancy packet over the unidirectional network.

2. A method as recited in claim 1, wherein each data packet has multiple bytes, and further comprising the step of generating the redundancy packet to include multiple bytes which correspond to the multiple bytes of the data packets, each byte in the redundancy packet containing redundancy information that is computationally derived from the corresponding bytes in each of the data packets within the redundancy group.

3. A method as recited in claim 1, wherein each data packet has multiple bytes, and further comprising the step of generating the redundancy packet to include multiple bytes which correspond to the multiple bytes of the data packets, each byte in the redundancy packet containing redundancy information that is computationally derived according to an XOR function of the corresponding bytes in each of the data packets within the redundancy group.

4. A method as recited in claim 1, further comprising the step of assigning a redundancy group identifier to every data packet within the redundancy group to identify the data packets as belonging to the redundancy group.

5. A method as recited in claim 1, further comprising the step of assigning a unique packet identifier to each data packet within the redundancy group.

6. A method as recited in claim 1, further comprising the step of assigning a content identifier to each data packet to identify the data packet as belonging to a particular data transmission.

7. A method as recited in claim 1, further comprising the following steps:

receiving the data packets and redundancy packet;
reconstructing any data packet from the redundancy group using the received data packets and the redundancy packet.

8. A method as recited in claim 1, further comprising the following steps:

forming multiple redundancy groups, each redundancy group having multiple data packets and at least one redundancy packet;

interleaving data packets from different ones of the redundancy groups prior to sending the data packets over the unidirectional network.

9. A computer-readable medium having computer-executable instructions for performing the steps of the method as recited in claim 1.

10. A computer programmed to perform the steps of the method as recited in claim 1.

11. A method for serving multiple data packets for a data transmission over a unidirectional network, comprising the following steps:

assigning a content identifier to each data packet to identify the data packet as belonging to the data transmission;

assigning a unique packet identifier to each data packet;
grouping the data packets into multiple redundancy groups, each redundancy group having multiple data packets;

10

assigning a redundancy group identifier to every data packet within a particular redundancy group to identify the data packets as belonging to the particular redundancy group;

forming at least one redundancy packet for associated ones of the redundancy groups, the redundancy packet containing redundancy information generated from the data packets in the associated redundancy group; and

assigning a redundancy group count to each data packet within the particular redundancy group to represent a total number of packets within the particular redundancy group.

12. A method as recited in claim 11, wherein each data packet has multiple bytes, and further comprising the step of generating the redundancy packet for each redundancy group to include multiple bytes which correspond to the multiple bytes of the data packets in the redundancy group, each byte in the redundancy packet containing redundancy information that is computationally derived from the corresponding bytes in each of the data packets within the redundancy group.

13. A method as recited in claim 11, wherein each data packet has multiple bytes, and further comprising the step of generating the redundancy packet for each redundancy group to include multiple bytes which correspond to the multiple bytes of the data packets in the redundancy group, each byte in the redundancy packet containing redundancy information that is computationally derived according to an XOR function of the corresponding bytes in each of the data packets within the redundancy group.

14. A method as recited in claim 11, further comprising the step of sending the data packets and the redundancy packets over the unidirectional network.

15. A method as recited in claim 11, further comprising the following steps:

interleaving data packets from different ones of the redundancy groups; and

sending the interleaved data packets along with the redundancy packets over the unidirectional network.

16. A computer-readable medium having computer-executable instructions for performing the steps of the method as recited in claim 11.

17. A computer programmed to perform the steps of the method as recited in claim 11.

18. A method for handling data packets transmitted over a unidirectional network, comprising the following steps:

receiving data packets carried over the unidirectional network, the data packets belonging to various redundancy groups;

receiving redundancy packets carried over the unidirectional network, the redundancy packets belonging to the various redundancy groups and containing redundancy information generated from the data packets in associated redundancy groups;

gathering the data packets and the redundancy packets into respective ones of the redundancy groups;

checking each redundancy group to determine if any of the data packets are missing; and

in an event that a particular data packet is missing from a particular redundancy group, reconstructing the particular data packet from remaining ones of the data packets in the particular redundancy group and the redundancy packet for the particular redundancy group.

19. A method as recited in claim 18, wherein each data packet has a redundancy group identifier to identify the data packet as belonging to one of the redundancy groups, and

6,081,907

11

said gathering step comprises gathering the data packets with like redundancy group identifiers to form the redundancy groups.

20. A method as recited in claim 18, wherein each data packet has a redundancy group count to represent a total number of packets contained within the redundancy group to which the data packet belongs, and said checking step comprises the step of detecting whether any data packet is missing from the particular redundancy group by comparing the redundancy group count in one of the data packets within the particular redundancy group with an actual number of packets gathered for the particular redundancy group.

21. A method as recited in claim 18, wherein each data packet has a unique packet identifier, and further comprising the step of determining which data packet is missing from the particular redundancy group.

22. A method as recited in claim 18, wherein each data packet and each redundancy packet has multiple bytes arranged in predefined positions so that byte positions within the redundancy packet correspond to byte positions within the data packet, and said reconstructing step comprises computationally deriving bytes in the particular data packet that is missing from the corresponding bytes in each of the remaining data packets and the redundancy packet within the particular redundancy group.

23. A method as recited in claim 18, wherein each data packet and each redundancy packet has multiple bytes arranged in predefined locations so that byte positions within the redundancy packet correspond to byte positions within the data packet, and said reconstructing step comprises XORing the corresponding bytes in each of the remaining data packets and the redundancy packet within the particular redundancy group to regenerate the bytes in the particular data packet that is missing.

24. A computer-readable medium having computer-executable instructions for performing the steps of the method as recited in claim 18.

25. A computer programmed to perform the steps of the method as recited in claim 18.

26. In a data delivery system for sending data packets from a server to a client over a unidirectional network, the client and the server each having a computer-readable medium, the computer-readable media on the server and the client having computer-executable instructions for performing steps comprising:

- grouping multiple data packets into a redundancy group at the server;
- generating at least one redundancy packet containing redundancy information derived from the data packets in the redundancy group;
- sending the data packets and the redundancy packet over the unidirectional network; and
- reconstructing at the client any data packet missing from the redundancy group using remaining ones of the data packets in the redundancy group and the redundancy packet for the redundancy group.

27. Computer-readable media having computer-executable instructions as recited in claim 26, further comprising computer-executable instructions to perform steps of forming multiple redundancy groups, each redundancy group having multiple data packets and at least one redundancy packet, and interleaving data packets from different ones of the redundancy groups at the server prior to sending the data packets over the unidirectional network.

28. A data delivery system for sending data packets from a server to a client over a unidirectional network, comprising:

12

a redundancy formatter located at the server, the redundancy formatter grouping multiple data packets into a redundancy group and generating at least one redundancy packet containing redundancy information derived from the data packets in the redundancy group; and

a packet rebuilder located at the client, the packet rebuilder reconstructing any data packet missing from the redundancy group using remaining ones of the data packets in the redundancy group and the redundancy packet for the redundancy group.

29. A data delivery system as recited in claim 28, wherein each data packet has multiple bytes at predefined positions within the packet, the redundancy formatter generating the redundancy packet to include multiple bytes which correspond in position to the bytes of the data packets, each byte in the redundancy packet containing redundancy information that is computationally derived from the corresponding bytes in each of the data packets within the redundancy group.

30. A data delivery system as recited in claim 28, wherein each data packet has multiple bytes at predefined positions within the packet, the redundancy formatter generating the redundancy packet to include multiple bytes which correspond in position to the bytes of the data packets, each byte in the redundancy packet containing redundancy information that is computationally derived according to an XOR function of the corresponding bytes in each of the data packets within the redundancy group.

31. A data delivery system as recited in claim 28, wherein each data packet comprises a redundancy group identifier to identify the data packet as belonging to the redundancy group.

32. A data delivery system as recited in claim 28, wherein each data packet comprises a unique packet identifier.

33. A data delivery system as recited in claim 28, wherein each data packet comprises a content identifier to identify the data packet as belonging to a particular data transmission.

34. A data delivery system as recited in claim 28, wherein the redundancy formatter forms multiple redundancy groups, with each redundancy group having multiple data packets and at least one redundancy packet, the redundancy formatter interleaving data packets from different ones of the redundancy groups prior to transmission over the unidirectional network to the client.

35. A computer-readable memory having a packet structure for a data packet to be carried by a unidirectional network as part of a data transmission, the data packet belonging to a group of data packets, the network packet structure comprising:

- a header comprising:
 - a packet identifier field to uniquely identify the data packet;
 - a content identifier field to uniquely identify the data packet as belonging to a particular data transmission;
 - a redundancy group identifier field to identify the packet as belonging to a redundancy group of data packets, the redundancy group comprising multiple data packets and at least one redundancy packet derived from the data packets;
 - a redundancy group count to represent a total number of the packets contained within the redundancy group to which the data packet belongs; and
- a data payload joined to the header to hold the data of the data packet.

Development of unidirectional data diode system in the secure environment

A. G. Vorontsov
Information Technology
FSUE VNIIA
Moscow, Russia
voroncov-ag@narod.ru

S. A. Petunin
Information Technology
FSUE VNIIA
Moscow, Russia
S.A.Petunin@gmail.com

The paper describes features of the development of unidirectional network (DataDiode devices) in the secure environment (the development of the unidirectional data transfer from the wide area network (WAN) to the secured enterprise network). A brief overview of the existing devices is given and their characteristics are described. The method to address limitation issues, such as lack of feedback channel related to the implementation and integration of one-way channel is described.

The main part of the research is focused on enhancing accessing process to the restricted network. The model of the unidirectional data transfer system, including DataDiode device, has been introduced.

The paper reflects:

- General principles of unidirectional data transfer;
- Topological features of the unidirectional data transfer model;
- Interaction with information security system;
- Providing load balancing and interactivity under high-load conditions.

Keywords: *data diode; unidirectional; security; highload; fiber; one-way gateway*

I. INTRODUCTION

A unidirectional gateway is a network appliance allowing data to be transferred only in one direction [1]. It doesn't allow data to pass in the opposite direction and connects different segments of the network with various privacy levels of the data processing and storage [2]. In the subject area under consideration, such one-way network solution must ensure no data transmission capacity in the opposite direction at the hardware level to preclude reconfiguration of security policy or firewall rules.

The class of such devices that provides isolation of network segments to prevent unauthorized access to the network information assets is referred to as Data Diode. Thus, it ensures the required data input to the closed network, and at the same time, it prevents unauthorized outputting of restricted information or any other external accessing to the closed network [3].


Currently, the Data Diode devices may be topologically different unidirectional data gateway. Some types of such systems are discussed below.

We introduce the notion of a private network segment and public network segment. The private network segment is a protected automated network system. The public network is an unsecured wide area network.

Schematic [4] of the unidirectional data gateway is shown in Fig. 1.

Figure 1. Data transmission to the downstream network


This Data Diode system is based on physically isolated fiber-optic communication and data feedback principle (no



network autonegotiation). The functional diagram of the system kernel is shown in Fig. 2, which demonstrates the absence of a data feedback.

Figure 2. Functional diagram of transmitting system kernel

The devices based on the given functional diagrams have a



number of advantages and disadvantages. Using Data Diodes allows resolving the problem of open data input to the downstream network, and implementing asynchronous management (without data feedback) of services within downstream network by means of control commands sent from the public network segment. It allows moving away from the previously used input methods (e.g. specialized input points where all the data is controlled by an operator) and to partially automating data input process.

However, the absence of data feedback brings forth a number of possible technical challenges. Implementation of the unidirectional channel causes the problem of transferred data verification which in the most cases is critical for data integrity:

- TCP/IP protocol requires handshaking;
- Data feedback is required for determination and adjustment of the communication rate;
- It is impossible to run application software since it requires a data feedback or data received acknowledgement.
- Most of web-services are not functional.

There are different approaches to resolve such problems [5]. For instance, the USA patent No. 5703562 Method for transferring data from an unsecured computer to a secured computer proposes a verification mechanism that uses a warning device coupled to a secured computer and emits a warning signal if an error was introduced during data transmission. The patent suggests using a single long duration tone [6]. It is obvious that the suggestion doesn't allow transmitting the checksum calculations from a Send Node to a Receive Node to allow the latter match the results, and thus define received data integrity.

Also, there are simplified approaches that deploy commercially-available hardware devices for establishment of a topologically adjusted network, configured only to one-way data transmission. However, these methods and approaches don't provide the sufficient protection level and it doesn't protect from attacks aimed to reverse data channel flow.

Out of the existing and functional devices, the most interesting configurations are built on the basis of the diagram shown in Fig. 2.

A. Unidirectional gateways with two module and repository

It consists of two components – the transmitter of the public network and the receiver of the private network coupled without any real data feedback. Every module contains an SSD drive with XFS file system. Both components are controlled by OS Linux special build.

The following operation algorithm is used [7]:

1. The public network data flows to the receiver, and then is written to an SSD drive.
2. As data transmission to the receiver is over, the data is synchronized with the private network receiver's SSD drive via UDP protocol that doesn't require acknowledgement, thus partially eliminates the necessity of data feedback. During synchronization process, the transmitted data is splitting into data packets, where each packet entered the checksum.
3. The data stored on the receiver, becomes available to a recipient, while the data on the public network SSD drive receiver is automatically deleted after synchronization is over.

Schematic diagram of two-module unidirectional gateways is shown in Fig. 3.




Figure 3. Functional diagram for two-module unidirectional gateways

Such implementation allows for one-way data transfer without data feedback required to confirm the integrity of the transmitted data and dynamic bandwidth adjustment of transfer rate. The data integrity is confirmed by the internal software. As the channel bandwidth and specification for the transmitter and the receiver is beforehand known, this provides an opportunity to determine synchronization speed in advance between SSDs drives.

The transmission speed has a great relevance because if the asynchronous behavior occurs, the receiver would not be able to keep up handling the input data stream, and the sender (the upstream network transmitter) would not know about it. This situation may lead to transmission failure.

The implementation disadvantages:

- XFS file system instability;
- Poor performance (about 30 Mb/s)
- The limited number of users is supported (100 users in total, 20 active users)
- The volume of SSD drives within the system is limited.

The most important measures of the device performance entered as the parameters listed in Table 1.

TABLE I. DATA TRANSFER RESULTS

Quantity*volume Measures	1x8000 Mb	3x8000 Mb	1030x7 Mb	3030x10 Mb
Average send speed, Mb/s	20	20	15	14
Average sync speed, Mb/s	30	30	25	25
Average receive speed, Mb/s	25	25	23	22
Average transmit time, min	18	53	20	82
Average number of transmission errors	0	0	15	43
Average number of sync errors	0	0	20	1315
Number of transmissions	5	5	5	5

B. End-to-end unidirectional gateway with one module

Unlike the previously two-module architecture, it is a single device. It is also a transmission device based on the same principles as the two-module unidirectional gateways but it doesn't have any internal storage and embedded software to perform data synchronization.

Operating principle is also build on the transmitting system kernel as the two-module device, though it has its specific differences [8]:

- Absence of the internal storage;
- Supports large number of active users (511);
- High performance (up to 90 Mb/sec).

Since there is no need to store the data within the device and no sync role – the given architecture is more fault tolerant and much better to implement. The software which responsible for data receiving and transmitting is located within the device on separate servers. The diagram of the device is shown in Fig. 4.

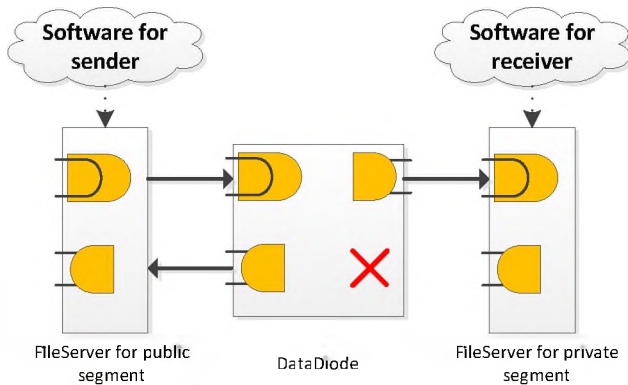


Figure 4. Scheme of the end-to-end unidirectional gateway

The most important measures of the device performance entered as the parameters listed in Table 2.

TABLE II. DATA TRANSFER RESULTS

Quantity * volume Measures	1x8000 Mb	3x8000 Mb	1030x 7 Mb	3030x 10 Mb
Average send speed, Mb/s	87	92	52	50
Average receive speed, Mb/s	87	92	52	50
Average transmit time, min	95	265	140	610
Average number of sync errors	0	0	7	0
Average number of redundancy data packets	1	1	1	2
Number of transmissions	5	5	5	5

II. UNIDIRECTIONAL DATA TRANSFER SYSTEM REQUIREMENTS

Necessity of development methods and a number of means to implement the automated unidirectional network ensure:

- guaranteed unidirectional data transfer;
- absence of transmitted data loss;
- communication channel performance (no less than one Gb/s);
- dynamic load balancing between output points from the public network;
- dynamic load balancing between input point to the private network;
- compatibility with information protection software;
- private network topology hiding;
- integration with network domain structure.

The proposed methods and means are automated, integrated with public and private network and has centralized management. The necessary condition to be able to interact with private network is the appropriate certification of Data Diode device.

III. SYSTEM MODEL

To develop the system model, it is necessary to define the objects.

Objects of the public network:

- A transmitting file server;
- Domain controller;
- User workstations.

Objects of the private network:

- A receiving file server;
- General file server;
- Domain controller;
- Network firewall;
- Data Diode;
- User workstations.

The proposed model implements network firewall in compliance with different network classes interaction requirements.

The functional diagram is shown in Fig. 5.

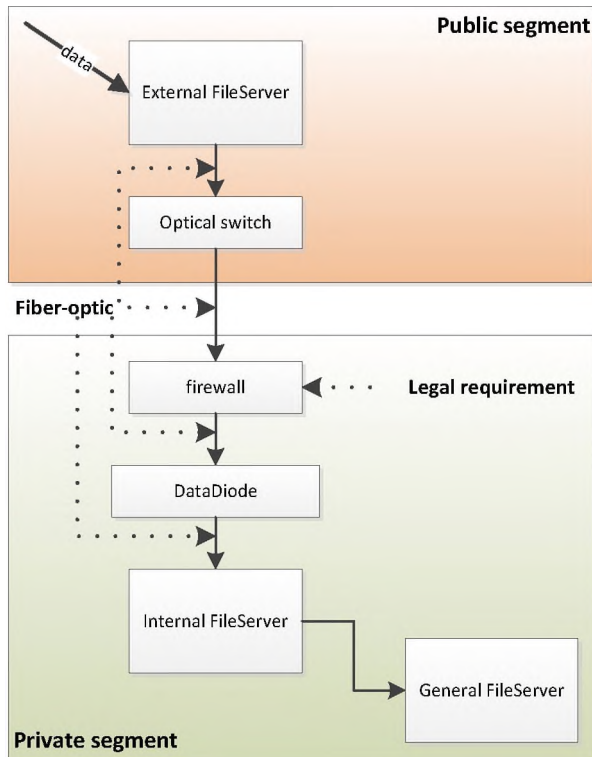


Figure 5. Example of a figure caption. (figure caption)

IV. STRUCTURAL IMPLEMENTATION METHODS

Architecture of the proposed solution is shown above. Besides customizing of communication equipment and sync software for the transmitting and receiving file servers, a number of problems need to be solved, such as:

- The incoming file queue problem;
- Overload of the devices that may become a reason of out of sync behavior;
- Interaction with information security system;
- Load balance;
- Interactive management and automation of the research system.

The problem of the incoming file queue is the selection of file processing priority. Before transmitting the current file in queue, the analysis of the whole ready-to-transmit queue is required, arranged by the file size, for example. Then, it is necessary to rebuild transmitting queue. This action solves such problems as large packet transmission and concentration

of a great number of small-sized files in file queue. A queue manager is implemented into the proposed model. If the large-sized file is being in the transmitting process while the system has only one unidirectional channel, it is necessary to wait until the file transfer is over. To resolve this problem file queue manager should compress large-sized files into equally sized archive volumes as part of the file transfer preparation procedure. It will enhance opportunities for queue management.

As it was mentioned above, absence of data feedback and transmission error that may occur due to asynchronous behavior of data transfer speed has two potential options:

- Application of sync software;
- Application of additional control facility for received files.

For the purpose of reduction such risk it is necessary to have identical technical characteristics for transmitting and receiving servers, and also contain high-performance disk subsystems, which preferably should be based on SSD drives.

The above mentioned method for file integrity check (uses application responsible for data transfer) needs optimization. While using this method, a possible problem may occur when it comes to huge number of files transfer that becomes more complicated regarding to inability to know the size, quantity and list of transmitted files due to a lack of data feedback.

The additional check method involves full data compressing and creation of archive volumes. It guarantees visual check for complete transfer of every archive file. If error occurs, the archive would not be upload to the receiving file server, thus a user would have to retransmit the archive volume.

Interaction with information protection system when end-to-end unidirectional gateway is used becomes a trivial task, aimed to ensure integration of transmitting and receiving file servers into the existing systems [9]. The main task is to provide error-free interaction with synchronizing software without any delay of network protocol or files reading and writing operations.

In case of using one unidirectional gateway device load balance is applied if exceeding or insufficient computational resources. To resolve such problem the scaling method is applied. It involves increasing the number of transmitting and receiving file servers and unidirectional channels. It is also necessary to apply balancing gateway that redirects the user to one of the receiving servers of the public network [10]. This architecture is shown in Fig. 6.

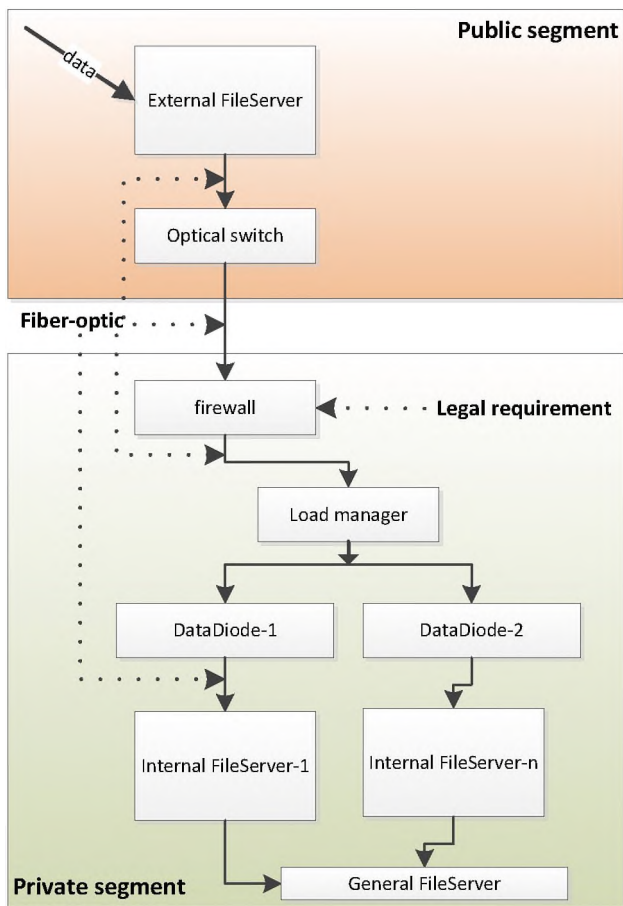


Figure 6. Example of a figure caption. (figure caption)

Development of a general control structure to ensure interactive control of the system is also required.

It provides:

- Addition of new network users;
- Automated control for transferred and received files that include archiving, transmission preparation, priority queuing, receiving files allocation.

The structure functions in cooperation with synchronizing software of the unidirectional gateway within public and private network. The conditions of no data feedback are observed.

Work algorithm of control structure and internal synchronization software within public network is shown in Fig. 7.

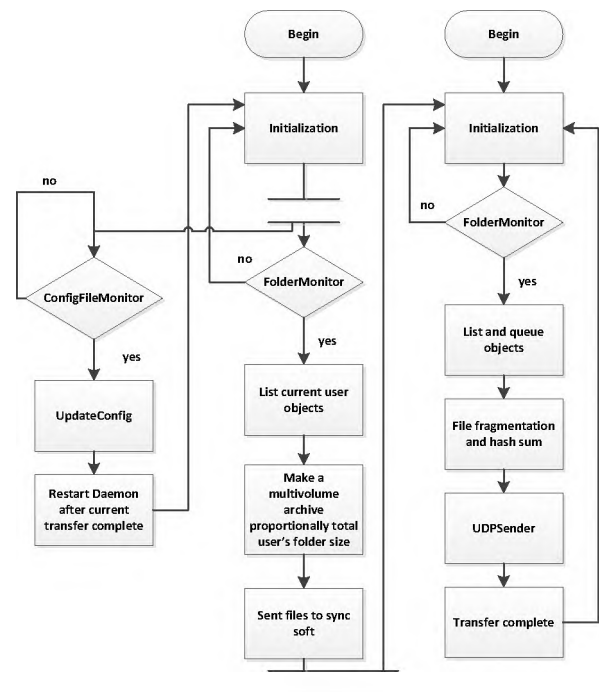


Figure 7. Sync software and control structure of the upstream network flowchart

Work algorithm of receiving data within private network is shown in Fig. 8.

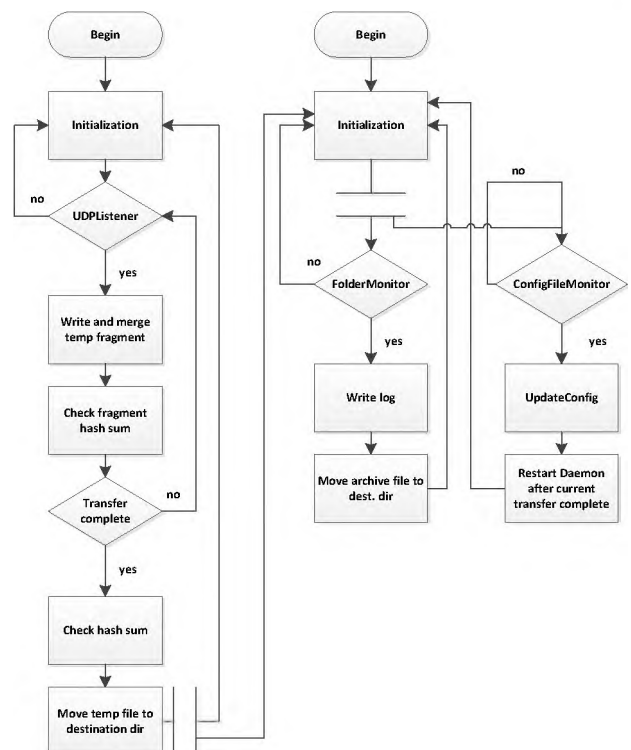


Figure 8. Sync software and control structure of the private network flowchart

V. CONCLUSION

This paper introduces and describes the development of hardware and software solutions for unidirectional gateway implementation in the conditions of secured environment. The actual implementation provided a unidirectional data flow from Wide Area Network to Local Area Network under the given conditions.

A primary result of this paper is theoretical development and implementation of unidirectional network system integrated into corporative network ensuring interaction with security systems. The developed control structure permits to scale the system considering basic load factors such as transmitted data volume and the number of users in system.

REFERENCES

- [1] Unidirectionalnetwork:
https://en.wikipedia.org/wiki/Unidirectional_network
- [2] Australian Government Information Management Office 2003, Securing systems with Starlight, Department of Finance and Administration http://www.agimo.gov.au/archive/publications_noie/2003/06/transform/defence.html
- [3] “CanSec” company – “What is unidirectional gateway”: <http://cansec.ru/21/unidirectional-gateway.html>
- [4] AMT Group – “InfoDiode - Unidirectional gateway system”: http://www.amt.ru/rubr.aspx?rubr_id=237&art_id=1154
- [5] Douglas W. Jones and Tom C. Bowersox “Secure Data Export and Auditing using Data Diodes” // 2006 USENIX/ACCURATE Electronic Voting Technology Workshop, 1 August 2006, Vancouver, https://www.usenix.org/legacy/events/evt06/tech/full_papers/jones/jones_html/
- [6] OKB SAPR – “Organization of a unidirectional data transmission channel on the basis of a protected service information carrier”: http://www.okbsapr.ru/lydin_tezisy2013_1.html
- [7] “CanSec” company – “Integrated file and mail unidirectional gateway Strom-File”: http://cansec.ru/products/strom_file.html
- [8] “CanSec” company – “High speed unidirectional gateway Strom-1000”: <http://cansec.ru/products/strom-1000.html>
- [9] Vorontsov A. G., Petunin S. A., Konyshev A. V. “Empowering information security systems in the conditions of domain environment” // Workshop on computer science and information technologies CSIT’2015, Rome, Italy, 2015.
- [10] Somerdata – “AROW Data Diode”: http://somerdata.com/?page_id=176

Data Diodes in Support of Trustworthy Cyber Infrastructure

Hamed Okhravi^{*}
University of Illinois at Urbana-Champaign
1308 West Main St.
Urbana, IL
okhravi@mit.edu

Fredrick T. Sheldon
Oak Ridge National Laboratory (ORNL)
One Bethel Valley Rd, M/S 6418
Oak Ridge, TN
sheldonft@ornl.gov

ABSTRACT

Interconnections between process control networks and enterprise networks has resulted in the proliferation of standard communication protocols in industrial control systems which exposes instrumentation, control systems, and the critical infrastructure components they operate to a variety of cyber attacks. Various standards and technologies have been proposed to protect industrial control systems against cyber attacks and to provide them with confidentiality, integrity, and availability. Among these technologies, data diodes provide protection of critical systems by the means of physically enforcing traffic direction on the network. In order to deploy data diodes effectively, it is imperative to understand the protection they provide, the protection they do not provide, their limitations, and their place in the larger security infrastructure. In this work, we briefly review the security challenges in an industrial control system, study data diodes, their functionalities and limitations, and propose a scheme for their effective deployment in trusted process control networks (TPCNs.)

Categories and Subject Descriptors

C.2.0 [Computer-Communication Networks]: General—*Security and protection*; B.4.1 [Input/Output and Data Communication]: Data Communications Devices

General Terms

Security

Keywords

Data Diodes, Trusted Process Control Networks, Industrial Control Systems

1. OVERVIEW OF PCNS AND SECURITY CHALLENGES

Figure 1 illustrates a typical process control network (PCN) architecture with paired firewall. In this architecture, the

^{*}The author is currently with MIT Lincoln Laboratory.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. CSIIRW '10, April 21-23, Oak Ridge, Tennessee, USA Copyright 2010 ACM 978-1-4503-0017-9 ... \$5.00

PCN contains the low level control devices such as programmable logic controllers (PLCs), remote terminal units (RTUs), master terminal unit (MTU), and the operator console. The enterprise network often contains the workstations and high level management consoles. The data historian sits in the demilitarized zone (DMZ) of the firewalls and acts as an intermediary between the PCN and EN. In fact, to protect the PCN from attacks and breaches going through the EN, status data is collected from the historian and not from the PCN directly.

Protecting PCNs often faces several challenges. Firewall configuration errors may result in unwanted traffic going to the PCN or legitimate traffic being dropped. In fact a study by Wool [15] shows that 80% of firewall rule sets allow any service on inbound traffic and insecure access to firewalls. Moreover, a firewall maybe bypassed by an attacker using encrypted tunnels (e.g. VPN) or unsecured out-of bound communication (e.g. dial-up maintenance connection.) Vulnerable end devices also pose a threat to the security of PCNs. Software/configuration bugs in the control devices may be exploited by an attacker to gain illegitimate access to the system or change the configuration of the critical components. Unsecured physical access to any part of the network (unsecured Ethernet ports) may also result in a benign or malicious damage to the PCNs. In addition, untrusted (rogue) devices or users may enter the network and breach its security. Finally, all of the above mentioned mechanisms may introduce malware (worms and viruses) to the critical systems.

2. DATA DIODES

Data diodes provide a physical mechanism for enforcing strict unidirectional communication between two networks. They are often implemented by removing transmitting component from one side and receiving component from another side of a bidirectional communication system (e.g. a fiber optic system with TX capability in only one side and RX capability in the other side.) Data diodes can only send information from one network (a.k.a the “low” network) to another network (a.k.a the “high” network.) The high network often contains data with higher classification level than the low network. Figure 2 illustrates two networks connected by a data diode.

2.1 Protection Provided

Data diodes can provide strong confidentiality from the high network to the low network; i.e. provided that the unidirectional

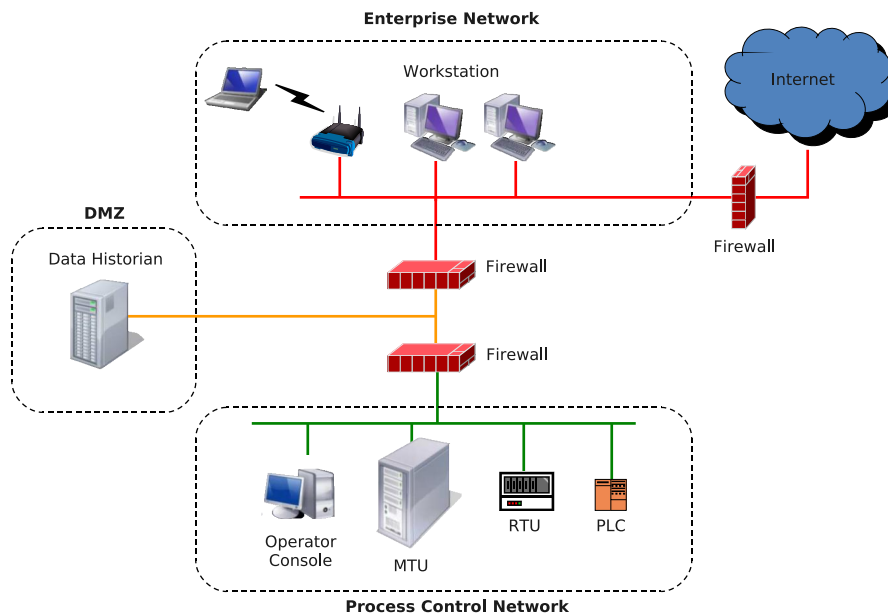


Figure 1: A typical paired-firewall industrial control system.




Figure 2: Two networks connected by a data diode.

tional connection is the only communication link between these two networks, information can flow from low to high, but there is no backflow of data. In a dual fashion, data diodes can provide strong integrity from the low network to the high network; i.e. a malicious component in the high network cannot corrupt data or perform network-based attacks on the low network (availability).

2.2 Protection Not Provided

It is sometimes claimed that data diodes protect the high network against cyber attacks. This, in fact, is not correct. Many cyber exploits do not require a session or bidirectional communication. Often fast propagating worms or malware need just one packet of data to infect a machine. Self expanding malware or quine programs [7] even limits the number of bytes required in the packet [13].

Moreover, in industrial control systems, the process control network is the critical component of the system for which availability and integrity are important properties. If the process control network is connected to the “high” side, the data diode does not protect it against breaches from the low network.

2.3 Limitations

A major limitation of the data diode is that it does not work with the standard TCP/IP protocols. It needs proprietary unidirectional protocols that do not require acknowledgments. On both sides of a data diode, gateways translate unidirectional protocols to standard bidirectional protocols to connect the diode to the rest of the network [4]. However, more high-end products [3] also accept TCP or UDP packets as input. Data diodes can be used to enhance security, but they are by no means even a nearly complete solution. They have to be placed carefully in conjunction with other defensive mechanisms.

2.4 Implementation

Data diodes are often implemented using serial links (RS-232) or optical fiber. In the serial link implementation, one of the two data cables (from high to low) is removed. In optical data diodes, the transmitter of the high network and the receiver of the low network are removed.

A major disadvantage of the RS-232 implementation is that in addition to data lines, there are control lines defined in the standard along which data can potentially flow back to the low network. Hence, optical fiber is the preferred implementation of data diodes.

3. TRUSTED PROCESS CONTROL NETWORK WITH DATA DIODES

A TPCN architecture [12] deploys trusted network (TN) [2, 1] technology to establish trust in industrial control systems. It uses information about the hardware and software states of devices in admission and access control decisions. When a device first joins the network, its hardware and software are checked; based on these checks, the appropriate access control rules are applied dynamically to the user, device and traffic. A TPCN architecture uses existing standards, protocols, and hardware devices to extend the concept of “trust” to the network architecture.

A TPCN has the following components:

- Client device: Every client device must be evaluated prior to admission to a TPCN.
- Network Access Device (NAD): All connectivity to a TPCN is implemented via a NAD, which enforces policy. NAD functionality may exist in devices such as switches, routers, VPN concentrators and wireless access points.
- Authentication, Authorization, and Access Control (AAA) Server: maintains the policy and provides rules to NADs based on the results of authentication and posture validation.
- Posture Validation Servers (PVSs): evaluate the compliance of a client before it can join a TPCN. A PVS is typically a specialization for one client attribute (e.g., operating system version and patch or virus signature release).
- Posture Remediation Servers: provide remediation options to a client device in the case of non-compliance.
- Directory Server: authenticates client devices based on their identities or roles.
- Other Servers: These include trusted versions of Audit, DNS, DHCP and VPN servers.

A TPCN architecture is presented in Figure 3. A client device intending to join the network communicates its request to the NAD. The NAD establishes the client device’s identity using EAP over the 802.1x protocol and sends the results to the AAA server using the RADIUS protocol. The AAA server returns a list of posture validation requirements and the addresses of the appropriate PVSs. The client then validates its posture with each of the PVSs. If the client is in compliance, the results are sent to the AAA server using the HCAP protocol. On the other hand, if the client lacks one or more requirements, the appropriate posture remediation servers suggest remediation actions to the client. The directory server determines the client’s group or role. Given all the results from the PVSs and the directory server, the AAA server determines the set of rules that apply to the client’s access and traffic and sends them to the NAD for enforcement.

From this point on, the client is permitted to communicate via the NAD and all its activities are monitored for policy

compliance. The policy held by the AAA server is in the form of an authentication requirement and a list of posture validation requirements.

When a client device joins the network, a NAD communicates with an AAA server on behalf of the device. The AAA server authenticates the device and provides rules based on the device’s security postures to the NAD. From this point on, the NAD enforces the policy on all ingress and egress traffic to/from the device. For example, an RTU with valid firmware is allowed to communicate with the historian; all other traffic is blocked. Okhravi and Nicol [12] provide two examples to further clarify the workings of a TPCN. They also describe methods to enhance availability of TPCNs and limit the number of configuration errors.

A TPCN addresses many of the security challenges by providing defense-in-depth and extending trust to the process control devices [11]. TPCNs build a security infrastructure for mission critical systems. Data diodes can be used to enhance TPCN protection by strictly limiting traffic at some sensitive points.

An important component of the TPCN network that can benefit from data diodes and tolerate their limitations is the data historian. The firewalls are often configured to drop any traffic going from the data historian to the PCN. If a data diode is placed between the historian and the PCN, the critical control devices can still push their status data to the DMZ while no traffic can flow back. Another diode may also be placed between the DMZ and EN to protect the integrity of the historian. Note that in both cases the “high” end of the diode is connected to the less critical components. This protects the PCN against attacks from EN or DMZ, granting integrity and availability. The confidentiality of the data sent to historian is arguably less important than protecting the PCN.

4. RELATED WORK

Kang, et al. [9] first designed and implemented a network device, *network pump*, for limiting convert back flow of data across the network. Network pump keeps the communication bidirectional, but it queues and sends the acknowledgments at probabilistic times. Stevens and Pope [10] discuss different implementations of data diodes and their assurance levels and limitations. Jones and Bowersox [8] propose the use of data diodes in secure data exports for voting systems. Finally, Roach [14] demonstrates the application of data diodes in aircraft instrumentation systems. To the best of our knowledge, we are the first to propose the application of data diodes in industrial control systems and develop a security infrastructure based on TPCNs for effective deployment of data diodes in process control applications.

5. CONCLUSION AND FUTURE WORK

Data diodes can offer some protection in the expense of imposing some limitations to the system. To effectively deploy them in a system, it is important to fully understand their functionalities and limitations. They cannot offer a comprehensive security solution, yet they can enhance the security of the system if used with care. A TPCN presents a trusted infrastructure for industrial control systems that can remedy many of the security issues. Data diodes can be placed

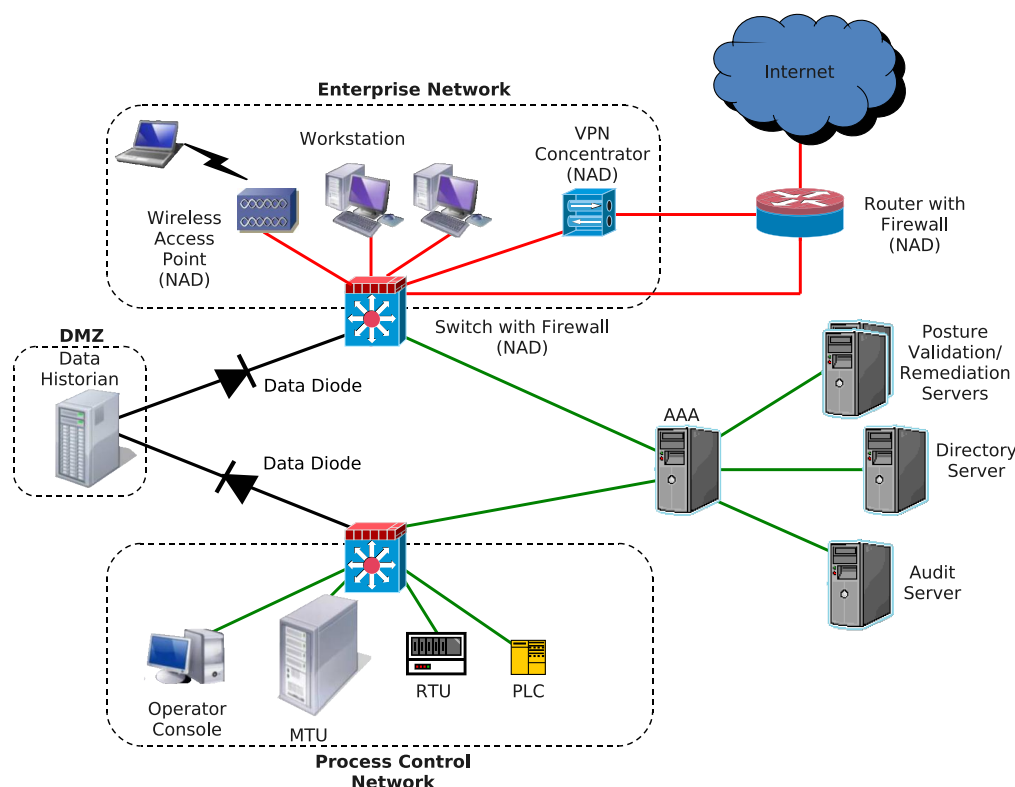


Figure 3: A TPCN with data diodes.

in sensitive places in a TPCN to protect the integrity of the control components and enhance the availability of the system. Based on our work on NAD rule conflicts [12], we plan to develop an algorithm to distribute firewall rules in the presence of data diodes in order to minimize rule conflicts [6] and implement a prototype on top of our testbed [5].

6. REFERENCES

- [1] Network Admission Control (NAC). Technical overview, Cisco Systems, Inc., 2005.
- [2] Getting started with Cisco NAC network modules in Cisco access routers. Technical manual, Cisco Systems, Inc., 2007.
- [3] Interactive Link Data Diode Device. Manual, BAE Systems, 2010.
- [4] Waterfall's Unidirectional Security Gateways. Manual, Waterfall, 2010. <http://www.waterfallsecurity.com/technology/>.
- [5] C. M. Davis, J. E. Tate, H. Okhravi, C. Grier, T. J. Overbye, and D. Nicol. SCADA cyber security testbed development. In *Proceedings of the 38th North American Power Symposium (NAPS 2006)*, pages 483–488, 2006.
- [6] A. Hari, S. Suri, and G. Parulkar. Detecting and resolving packet filter conflicts. In *Proceedings of IEEE INFOCOM*, pages 1203–1212, 2000.
- [7] D. R. Hofstadter. *Godel, Escher, Bach: An Eternal Golden*. Basic Books, Inc., New York, NY, 1 edition, 1979.
- [8] D. W. Jones and T. C. Bowersox. Secure data export and auditing using data diodes. In *EVT'06: Proceedings of the USENIX Electronic Voting Technology Workshop 2006*, pages 4–4, Berkeley, CA, USA, 2006. USENIX Association.
- [9] M. H. Kang, I. S. Moskowitz, and S. Chincheck. The pump: A decade of covert fun. In *ACSAC '05: Proceedings of the 21st Annual Computer Security Applications Conference*, pages 352–360, Washington, DC, USA, 2005. IEEE Computer Society.
- [10] S. M. and P. M. Data Diodes. Technical report - dsto-tr-0209, Electronics and Surveillance Research Laboratory (DSTO), 1995.
- [11] H. Okhravi and D. Nicol. Applying trusted network technology to process control systems. In E. Goetz and S. Sheno, editors, *Critical Infrastructure Protection II*, pages 57–70. Springer, Boston, MA, 2 edition, 2008.
- [12] H. Okhravi and D. Nicol. Application of trusted network technology to industrial control networks. *Elsevier International Journal of Critical Infrastructure Protection (IJCIP)*, 2(3):84–94, 2009.
- [13] M. R. Rieback, B. Crispo, and A. S. Tanenbaum. Is your cat infected with a computer virus? In *Proceedings of the Fourth Annual IEEE International Conference on Pervasive Computing and Communications*, pages 169–179, 2006.
- [14] J. Roach. The architecture of aircraft instrumentation networks. In *Proceedings of the International Telemetering Conference (ITC 2007)*, 2007.
- [15] A. Wool. A quantitative study of firewall configuration errors. *Computer*, 37(6):62–67, 2004.

EXHIBIT H

WEBSTER'S NEW WORLDTM

DICTIONARY *of* COMPUTER TERMS

EIGHTH EDITION

The Best Computer Dictionary in Print

Completely revised and updated

Contains extensive Internet coverage

More than 4,000 entries, terms, and acronyms

Bryan Pfaffenberger

THE NAME YOU TRUST

Dedication

For Suzanne, always

Webster's New World™ Dictionary of Computer Terms,
8th Edition

Copyright © 2000 by
IDG Books Worldwide, Inc.
An International Data Group Company
919 E. Hillsdale Blvd.
Suite 400
Foster City, CA 94404

All rights reserved including the right of reproduction in whole
or in part in any form.

For general information on IDG Books Worldwide's books in
the U.S., please call our Consumer Customer Service depart-
ment at 1-800-762-2974. For reseller information, including
discounts, bulk sales, customized editions, and premium sales,
please call our Reseller Customer Service department at
1-800-434-3422.

A Webster's New World™ Book

WEBSTER'S NEW WORLD DICTIONARY is a registered
trademark of IDG Books Worldwide, Inc.

Library of Congress Catalog Number: 98-68180

ISBN: 0-02-863777-1

Manufactured in the United States of America

1 2 3 4 5 6 7 00 1 02 03 04

laddered text In Unix and Unix-like operating systems, a printing defect caused by the printer's failure to start a new line after executing a line feed. This problem is caused by a basic incompatibility between DOS and Unix text files: DOS text files end with a carriage return and a line feed character, while Unix text files end with a line feed character only. The system administrator can fix this problem by altering a printer configuration file. See *carriage return* and *line feed*.

LAN Acronym for local area network. A computer network that physically links two or more computers within a geographically limited area (generally one building or a group of buildings). The linked computers are called workstations. Peer-to-peer LANs enable the linked computers to share expensive peripherals such as laser printers; client/server networks use a LAN server to make resources (such as databases and applications) available to workstation users. Local area networks have a characteristic topology (such as bus, ring, or star) and implement

312 **LAN-aware program**

one or more networking protocols (such as AppleTalk, Ethernet, or TCP/IP). See *AppleTalk*, *baseband*, *broadband*, *bus network*, *client/server*, *Ethernet*, *multiuser system*, *NetWare*, *network operating system (NOS)*, *peer-to-peer network*, *ring network*, and *star network*.

LAN-aware program A version of an application program specifically modified so that the program can function in a local area network (LAN) environment. Network versions of transactional application programs (database management programs) create and maintain shared files. An invoice-processing program, for example, has access to a database of accounts receivable. The network versions of nontransactional programs, such as word processing programs, include file security features to prevent unauthorized users from gaining access to your documents. LAN-aware programs boast features such as concurrency control, which manages multiple copies of files, and file locking and prevents unauthorized users from accessing certain files. LAN-aware programs usually are stored on a file server. See *LAN-ignorant program*.

LAN backup program A program designed specifically to back up the programs and data stored on a local area network's

EXHIBIT I

IEEE Std 802-2001®
(Revision of IEEE Std 802-1990™)

IEEE Standards

802®

IEEE Standard for Local and Metropolitan Area Networks: Overview and Architecture

IEEE Computer Society

Sponsored by the
LAN/MAN Standards Committee



Published by
The Institute of Electrical and Electronics Engineers, Inc.
3 Park Avenue, New York, NY 10016-5997, USA

7 February 2002

Print: SH94947
PDF: SS94947

IEEE Std 802®-2001 (R2007)
(Revision of IEEE Std 802-1990)

IEEE Standard for Local and Metropolitan Area Networks: Overview and Architecture

Sponsor

LAN/MAN Standards Committee
of the
IEEE Computer Society

Reaffirmed 21 March 2007
Approved 6 December 2001

IEEE-SA Standards Board

Abstract: IEEE Std 802-2001, IEEE Standards for Local and Metropolitan Area Networks: Overview and Architecture, provides an overview to the family of IEEE 802 Standards. It defines compliance with the family of IEEE 802 Standards; it describes the relationship of the IEEE 802 Standards to the Open Systems Interconnection Basic Reference Model [ISO/IEC 7498-1:1994] and explains the relationship of these standards to the higher layer protocols; it provides a standard for the structure of LAN MAC addresses; and it provides a standard for identification of public, private, and standard protocols.

Keywords: IEEE 802 standards compliance, Local Area Networks (LANs), LAN/MAN architecture, LAN/MAN reference model, Metropolitan Area Networks (MANs).

The Institute of Electrical and Electronics Engineers, Inc.
3 Park Avenue, New York, NY 10016-5997, USA

Copyright © 2002 by the Institute of Electrical and Electronics Engineers, Inc.
All rights reserved. Published 7 February 2002. Printed in the United States of America.

IEEE and 802 are registered trademarks in the U.S. Patent & Trademark Office, owned by the Institute of Electrical and Electronics Engineers, Incorporated.

Print: ISBN 0-7381-2940-2 SH94947
PDF: ISBN 0-7381-2941-0 SS94947

No part of this publication may be reproduced in any form, in an electronic retrieval system or otherwise, without the prior written permission of the publisher.

IEEE Standards documents are developed within the IEEE Societies and the Standards Coordinating Committees of the IEEE Standards Association (IEEE-SA) Standards Board. The IEEE develops its standards through a consensus development process, approved by the American National Standards Institute, which brings together volunteers representing varied viewpoints and interests to achieve the final product. Volunteers are not necessarily members of the Institute and serve without compensation. While the IEEE administers the process and establishes rules to promote fairness in the consensus development process, the IEEE does not independently evaluate, test, or verify the accuracy of any of the information contained in its standards.

Use of an IEEE Standard is wholly voluntary. The IEEE disclaims liability for any personal injury, property or other damage, of any nature whatsoever, whether special, indirect, consequential, or compensatory, directly or indirectly resulting from the publication, use of, or reliance upon this, or any other IEEE Standard document.

The IEEE does not warrant or represent the accuracy or content of the material contained herein, and expressly disclaims any express or implied warranty, including any implied warranty of merchantability or fitness for a specific purpose, or that the use of the material contained herein is free from patent infringement. IEEE Standards documents are supplied “**AS IS.**”

The existence of an IEEE Standard does not imply that there are no other ways to produce, test, measure, purchase, market, or provide other goods and services related to the scope of the IEEE Standard. Furthermore, the viewpoint expressed at the time a standard is approved and issued is subject to change brought about through developments in the state of the art and comments received from users of the standard. Every IEEE Standard is subjected to review at least every five years for revision or reaffirmation. When a document is more than five years old and has not been reaffirmed, it is reasonable to conclude that its contents, although still of some value, do not wholly reflect the present state of the art. Users are cautioned to check to determine that they have the latest edition of any IEEE Standard.

In publishing and making this document available, the IEEE is not suggesting or rendering professional or other services for, or on behalf of, any person or entity. Nor is the IEEE undertaking to perform any duty owed by any other person or entity to another. Any person utilizing this, and any other IEEE Standards document, should rely upon the advice of a competent professional in determining the exercise of reasonable care in any given circumstances.

Interpretations: Occasionally questions may arise regarding the meaning of portions of standards as they relate to specific applications. When the need for interpretations is brought to the attention of IEEE, the Institute will initiate action to prepare appropriate responses. Since IEEE Standards represent a consensus of concerned interests, it is important to ensure that any interpretation has also received the concurrence of a balance of interests. For this reason, IEEE and the members of its societies and Standards Coordinating Committees are not able to provide an instant response to interpretation requests except in those cases where the matter has previously received formal consideration.

Comments for revision of IEEE Standards are welcome from any interested party, regardless of membership affiliation with IEEE. Suggestions for changes in documents should be in the form of a proposed change of text, together with appropriate supporting comments. Comments on standards and requests for interpretations should be addressed to:

Secretary, IEEE-SA Standards Board
445 Hoes Lane
P.O. Box 1331
Piscataway, NJ 08855-1331
USA

<p>Note: Attention is called to the possibility that implementation of this standard may require use of subject matter covered by patent rights. By publication of this standard, no position is taken with respect to the existence or validity of any patent rights in connection therewith. The IEEE shall not be responsible for identifying patents for which a license may be required by an IEEE standard or for conducting inquiries into the legal validity or scope of those patents that are brought to its attention.</p>

The IEEE and its designees are the sole entities that may authorize the use of the IEEE-owned certification marks and/or trademarks to indicate compliance with the materials set forth herein.

Authorization to photocopy portions of any individual standard for internal or personal use is granted by the Institute of Electrical and Electronics Engineers, Inc., provided that the appropriate fee is paid to Copyright Clearance Center. To arrange for payment of licensing fee, please contact Copyright Clearance Center, Customer Service, 222 Rosewood Drive, Danvers, MA 01923 USA; +1 978 750 8400. Permission to photocopy portions of any individual standard for educational classroom use can also be obtained through the Copyright Clearance Center.

IEEE-SA Trademark Usage/Compliance Statement

Proper usage of the trademark **IEEE Std 802-2001**[®] is mandatory and is to be followed in all references of the Standard. The mark **IEEE**[®] is the registered trademark of the Institute of Electrical and Electronics Engineers, Inc., and must be used in bold type. It is to appear with the registered trademark symbol “®” the first time “**IEEE**” appears in the text. The use of “IEEE Std xxxx-200x” should include the trademark symbol “TM” (e.g., **IEEE Std xxxx-200x**TM) at least the first time it is used in text, unless the number of the standard is also trademark registered (e.g., **802**[®]), then the symbol “®” must be used.


It is not permissible to use the standard number alone or with “**IEEE**” to indicate conformance or compliance with the associated standard. The user of the Standard should contact the Manager, Standards Licensing and Contracts for information concerning issues regarding indicating product compliance with an **IEEE** standard. To represent that a product has been designed to meet an **IEEE** standard, it is permissible to state that “the product has been engineered, designed or manufactured with the intent to meet the requirements of **IEEE Std xxxx-200x**TM”. However, it is *not* permissible to state or refer to a product as “xxxx compliant,” “xxxx certified,” “**IEEE xxxx** conformant,” “**IEEE xxxx** certified,” or the like, unless the user has obtained a Certification License from the Manager, Standards Licensing and Contracts.

Introduction

(This introduction is not part of IEEE Std 802-2001[®], IEEE Standard for Local and Metropolitan Area Networks: Overview and Architecture.)

IEEE Std 802-2001[®] provides an overview to the family of IEEE 802[®] Standards. It defines compliance with the family of IEEE 802[®] Standards; it describes the relationship of the IEEE 802[®] Standards to the Open Systems Interconnection Basic Reference Model [ISO/IEC 7498-1: 1994] and explains the relationship of these standards to higher layer protocols; it provides a standard for the structure of LAN MAC addresses; and it provides a standard for the identification of public, private, and standard protocols.

This standard is part of a family of standards for local and metropolitan area networks. The relationship between the standard and other members of the family is shown below. (The numbers in the figure refer to IEEE standard numbers.)



This family of standards deals with the Physical and Data Link Layers as defined by the International Organization for Standardization (ISO) Open Systems Interconnection Basic Reference Model (ISO/IEC 7498-1:1994). The access standards define several types of medium access technologies and associated physical media, each appropriate for particular applications or system objectives. Other types are under investigation.

The standards defining the technologies noted above are as follows:

- IEEE Std 802[®]:¹ *Overview and Architecture.* This standard provides an overview to the family of IEEE 802[®] Standards. This document forms part of the IEEE Std 802.1[®] scope of work.
- IEEE Std 802.1B[®] and 802.1K[®] [ISO/IEC 15802-2]: *LAN/MAN Management.* Defines an Open Systems Interconnection (OSI) management-compatible architecture, and services and protocol elements for use in a LAN/MAN environment for performing remote management.

¹The IEEE 802[®] Architecture and Overview Specification, originally known as IEEE Std 802.1A[®], has been renumbered as IEEE Std 802[®]. This has been done to accommodate recognition of the base standard in a family of standards. References to IEEE Std 802.1A[®] should be considered as references to IEEE Std 802[®].

- IEEE Std 802.1D[®] *Media Access Control (MAC) Bridges*. Specifies an architecture and protocol for the [ISO/IEC 15802-3]: interconnection of IEEE 802[®] LANs below the MAC service boundary.
- IEEE Std 802.1E[®] [ISO/IEC 15802-4]: *System Load Protocol*. Specifies a set of services and protocol for those aspects of management concerned with the loading of systems on IEEE 802[®] LANs.
- IEEE Std 802.1F[®] *Common Definitions and Procedures for IEEE 802[®] Management Information*.
- IEEE Std 802.1G[®] [ISO/IEC 15802-5]: *Remote Media Access Control (MAC) Bridging*. Specifies extensions for the interconnection, using non-LAN systems communication technologies, of geographically separated IEEE 802[®] LANs below the level of the logical link control protocol.
- IEEE Std 802.1H[®] [ISO/IEC TR 11802-5] *Recommended Practice for Media Access Control (MAC) Bridging of Ethernet V2.0 in IEEE 802[®] Local Area Networks*.
- IEEE Std 802.1Q[®] *Virtual Bridged Local Area Networks*. Defines an architecture for Virtual Bridged LANs, the services provided in Virtual Bridged LANs, and the protocols and algorithms involved in the provision of those services.
- IEEE Std 802.2[®] [ISO/IEC 8802-2]: *Logical Link Control*.
- IEEE Std 802.3[®] [ISO/IEC 8802-3]: *CSMA/CD Access Method and Physical Layer Specifications*.
- IEEE Std 802.4[®] [ISO/IEC 8802-4]: *Token Bus Access Method and Physical Layer Specifications*.
- IEEE Std 802.5[®] [ISO/IEC 8802-5]: *Token Ring Access Method and Physical Layer Specifications*.
- IEEE Std 802.6[®] [ISO/IEC 8802-6]: *Distributed Queue Dual Bus Access Method and Physical Layer Specifications*.
- IEEE Std 802.10[®]: *Interoperable LAN/MAN Security*. Currently approved: Secure Data Exchange (SDE).
- IEEE Std 802.11[®]: [ISO/IEC 8802-11] *Wireless LAN Medium Access Control (MAC) Sublayer and Physical Layer Specifications*.
- IEEE Std 802.12[®]: [ISO/IEC 8802-12] *Demand Priority Access Method, Physical Layer and Repeater Specification*.
- IEEE Std 802.15[®]: *Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for: Wireless Personal Area Networks*.
- IEEE Std 802.16[®]: *Standard Air Interface for Fixed Broadband Wireless Access Systems*.
- IEEE Std 802.17[®]: *Resilient Packet Ring Access Method and Physical Layer Specifications*.

In addition to the family of standards, the following is a recommended practice for a common physical layer technology:

- IEEE Std 802.7[®]: *IEEE Recommended Practice for Broadband Local Area Networks.*

The reader of this standard is urged to become familiar with the complete family of standards.

Conformance test methodology

An additional standards series, identified by the number 1802[™], has been established to identify the conformance test methodology documents for the IEEE 802[®] family of standards. Thus the conformance test documents for IEEE 802.3[®] are numbered 1802.3[™], the conformance test documents for IEEE 802.5[®] will be 1802.5[™], and so on. Similarly, ISO will use 18802 to number conformance test standards for 8802 standards.

Participants

At the time this standard was completed, the IEEE 802.1[®] Working Group had the following membership:

William P. Lidinsky, Chair
Tony Jeffree, Vice Chair
Alan Chambers, Tony Jeffree, Editors

Steve Adams
 Floyd Backes
 John Bartlett
 Les Bell
 Michael Berger
 James S. Binder
 Bill Bunch
 Paul Carroll
 Jeffrey Catlin
 David W. Chang
 Hon Wah Chin
 Chris Christ
 Paul Congdon
 Glenn Connery
 David Cullerot
 Ted Davies
 David Delaney
 Jeffrey Dietz
 Peter Ecclesine
 J. J. Ekstrom
 Norman W. Finn
 Lars Henrik Frederiksen
 Anoop Ghanwani
 Sharam Hakimi
 John Hart
 Richard Hausman
 David Head
 Deepak Hegde

Ariel Hendel
 Steve Horowitz
 Altaf Hussain
 Vipin K. Jain
 Toyoyuki Kato
 Hal Keen
 Kevin Ketchum
 Keith Klamm
 Bruce Kling
 Dan Krent
 Paul Lachapelle
 Paul Langille
 Johann Lindmeyr
 Milan Merhar
 John Messenger
 Colin Mick
 Yaron Nachman
 Krishna Narayanaswamy
 Lawrence Ng
 Henry Ngai
 Satoshi Obara
 Toshio Ooka
 Jorg Ottensmeyer
 Luc Pariseau
 John Pickens
 Steve Ramberg
 Shlomo Reches
 James Richmond

Anil Rijasinghani
 Ayman Sayed
 Mick Seaman
 Lee Sendelbach
 Himanshu Shah
 Phil Simmons
 Paramjeet Singh
 Rosemary V. Slager
 Alexander Smith
 Andrew Smith
 Larry Stefani
 Stuart Soloway
 Sundar Subramaniam
 Robin Tasker
 Pat Thaler
 Steve Van Seters
 Dono van-Mierop
 John Wakerly
 Peter Wang
 Y. C. Wang
 Trevor Warwick
 Keith Willette
 Michael Witkowski
 Edward Wong
 Michael D. Wright
 Michele Wright
 Allen Yu
 Wayne Zakowski

The following members of the balloting committee voted on this standard. Balloters may have voted for approval, disapproval, or abstention.

Don Aelmore	Vic Hayes	Satoshi Obara
Jack S. Andresen	Donald N. Heirman	Charles Oestereicher
Lek Ariyavisitakul	Raj Jain	Roger Pandanda
Thomas W. Bailey	David V. James	Vikram Punj
Kevin Barry	Tony A. Jeffree	Sailesh Rao
Brad J. Booth	Henry D. Keen	Stanley A. Reible
Chris Byham	Peter M. Kelly	James A. Renfro
James T. Carlo	Stuart J. Kerry	Gary S. Robinson
David E. Carlson	Daniel R. Krent	Edouard Y. Rocher
Alan M. Chambers	Stephen Barton Kruger	James W. Romlein
Thomas J. Dineen	David J. Law	Floyd E. Ross
Christos Douligeris	Walter Levy	Christoph Ruland
Sourav K. Dutta	William P. Lidinsky	Robert Russell
Paul S. Eastman	Randolph S. Little	Norman Schneidewind
Richard Eckard	Joseph G. Maley	James E. Schuessler
Philip H. Enslow	Roger B. Marks	Gregory D. Schumacher
Changxin Fan	Peter Martini	Michael Seaman
John W. Fendrich	Richard McBride	Rich Seifert
Michael A. Fischer	Bennett Meyer	Leo Sintonen
Howard M. Frazier	Tremont Miao	Michael A. Smith
Ken J. Friedenbach	David S. Millman	Fred J. Strauss
Robert J. Gagliano	Mart Molle	Walter T. Thirion
Gautam Garai	John E. Montague	Geoffrey O. Thompson
Alireza Ghazizahedi	Masahiro Morikura	Mark-Rene Uchida
Patrick S. Gonia	Wayne D. Moyers	Scott A. Valcourt
Julio Gonzalez-Sanz	Shimon Muller	Paul A. Willis
Robert M. Grow	Paul Nikolich	Wayne Zakowski
David B. Gustavson	Erwin Noble	Johathan M. Zweig
Chris G. Guy	Ellis S. Nolley	
	Robert O'Hara	

When the IEEE-SA Standards Board approved this standard on 6 December 2001, it had the following membership:

Donald N. Heirman, *Chair*

James T. Carlo, *Vice Chair*

Judith Gorman, *Secretary*

Satish K. Aggarwal	James H. Gurney	James W. Moore
Mark D. Bowman	Richard J. Holleman	Robert F. Munzner
Gary R. Engmann	Lowell G. Johnson	Ronald C. Petersen
Harold E. Epstein	Robert J. Kennelly	Gerald H. Peterson
H. Landis Floyd	Joseph L. Koepfinger*	John B. Posey
Jay Forster*	Peter H. Lips	Gary S. Robinson
Howard M. Frazier	L. Bruce McClung	Akio Tojo
Ruben D. Garzon	Daleep C. Mohla	Donald W. Zipse

*Member Emeritus

Also included is the following nonvoting IEEE-SA Standards Board liaison:

Alan Cookson, *NIST Representative*

Donald R. Volzka, *TAB Representative*

Jennifer McClain Longman
IEEE Standards Project Editor

Contents

1.	Scope	1
1.1	General	1
1.2	Key concepts	1
1.3	Application and support	2
1.4	An international family of standards	3
2.	References	3
3.	Definitions, acronyms, and abbreviations	4
3.1	Definitions	4
4.	Abbreviations and acronyms	6
5.	Compliance	7
6.	Reference and implementation models	7
6.1	Introduction	7
6.2	RM description for end stations	10
6.3	Interconnection and interworking	12
7.	General requirements for an 802 [®] LAN or MAN	15
7.1	Services supported	15
7.2	Size and extent	16
7.3	Error rates	16
7.4	Service availability	16
7.5	Safety, and lightning and galvanic protection	16
7.6	Regulatory requirements	17
8.	LAN/MAN management	17
8.1	General-purpose LAN/MAN management	17
8.2	Special-purpose LAN/MAN management standards	19
9.	Universal addresses and protocol identifiers	19
9.1	OUI	20
9.2	48-bit universal LAN MAC addresses	20
9.3	Protocol identifiers	22
9.4	Standard group MAC addresses	23
9.5	Bit-ordering and different MACs	23
10.	Protocol discrimination above the MAC sublayer: Subnetwork Access Protocol (SNAP) and Ethernet types	27
10.1	Introduction	27
10.2	Basic concepts	27
10.3	Subnetwork Access Protocol	28

10.4 Ethernet types: Format, function, and administration	29
10.5 Encapsulation of Ethernet frames over LLC	29
11. ISLAN and MAN support for isochronous bearer services	30
11.1 Key concepts	30
11.2 Applications	31
11.3 Isochronous service access points and PhSAPs	31
11.4 ISLAN signaling	32
Annex A (informative) Bibliography (Additional references for LAN/MAN-related standards)	33

IEEE Standard for Local and Metropolitan Area Networks: Overview and Architecture

1. Scope

1.1 General

This document serves as the foundation for the family of IEEE 802[®] Standards published by IEEE for Local Area Networks (LANs) and Metropolitan Area Networks (MANs). It contains descriptions of the networks considered as well as a reference model (RM) for protocol standards. Compliance with the family of IEEE 802[®] Standards is defined, and a standard for the identification of public, private, and standard protocols is included.

1.2 Key concepts

The LANs described herein are distinguished from other types of data networks in that they are optimized for a moderate-sized geographic area, such as a single office building, a warehouse, or a campus. An IEEE 802[®] LAN is a peer-to-peer communication network that enables stations to communicate directly on a point-to-point, or point-to-multipoint, basis without requiring them to communicate with any intermediate switching nodes. LAN communication takes place at moderate-to-high data rates, and with short transit delays, on the order of a few milliseconds or less.

A LAN is generally owned, used, and operated by a single organization. This is in contrast to Wide Area Networks (WANs) that interconnect communication facilities in different parts of a country or are used as a public utility.

A MAN is optimized for a larger geographical area than is a LAN, ranging from several blocks of buildings to entire cities. As with local networks, MANs can also depend on communications channels of moderate-to-high data rates. A MAN might be owned and operated by a single organization, but it usually will be used by many individuals and organizations. MANs might also be owned and operated as public utilities. They will often provide means for internetworking of local networks.

Although primarily aimed at deployment on the scale of a large building or a campus, LANs are also frequently applied in smaller areas, such as small offices or single laboratories, and increasingly in homes. At the small-scale application level, a LAN is different from the type of network, such as a data bus or backplane bus, that is optimized for the interconnection of devices on a desktop or of components within a

single piece of equipment. However, desktop-scale applications of LANs are also possible, particularly where the nature of the application is more suited to peer-to-peer communication among autonomous components, as opposed to a system structure with more centralized control.

The original IEEE 802® LAN technologies used shared-medium communication, with information broadcast for all stations to receive. That approach has been varied and augmented subsequently, but in ways that preserve the appearance of simple peer-to-peer communications behavior for end stations. In particular, the use of bridges (see 6.3.2) for interconnecting LANs is now widespread. These devices allow the construction of networks with much larger numbers of LAN end stations, and much higher aggregate throughput, than would be achievable with a single shared-medium LAN. End stations attached to such a bridged LAN can communicate with each other just as though they were attached to a single shared-medium LAN (however, the ability to communicate with other stations can be limited by use of management facilities in the bridges, particularly where broadcast or multicast transmissions are involved). A further stage in this evolution has led to the use of point-to-point full duplex communication in LANs, either between an end station and a bridge or as a typically high-speed link between a pair of bridges.

The basic communications capabilities provided by all LANs and MANs are packet-based, as opposed to either cell-based or isochronous. That is, the basic unit of transmission is a sequence of data octets, which can be of any length within a range that is dependent on the type of LAN; for all LAN types, the maximum length is in excess of 1000 octets. (By contrast, cell-based communication transmits data in shorter, fixed-length units; isochronous communication transmits data as a steady stream of octets, or groups of octets, at equal time intervals.)

An optional function that may be offered by a LAN or a MAN is the provision of local networking of isochronous bearer services that are compatible with, or higher speed versions of, Integrated Services Digital Networks (ISDN) as defined by the ITU-T I-series Recommendations, to support voice, video, and data devices and terminals. These services are based on the use of end-user to end-user isochronous bearers that will span the supporting Integrated Services LAN (ISLAN) or MAN and an intervening ISDN-conformant WAN. Typically, the information streams for packet and isochronous services are multiplexed over the same physical media. In addition, capabilities are specified for a single integrated management of these various streams.

1.3 Application and support

The networks are intended to have wide applicability in many environments. The primary aim is to provide for moderate-cost devices and networks, suitable for commercial, educational, governmental, and industrial applications. Low-cost alternatives are possible for some networks, and application in other environments is not precluded. The following lists are intended to show some applications and devices and, as such, are not intended to be exhaustive, nor do they constitute a set of required items:

- File transfer
- Graphics
- Text processing
- Desktop publishing
- Electronic mail
- Database access
- Transaction processing
- Multimedia
- Office automation

- Process control
- Robotics
- Integrated Services (voice, video and data) applications
- Client/server applications

The networks are intended to support various data devices, such as the following:

- Computers
- Terminals
- Mass storage devices
- Printers and plotters
- Photocopiers and facsimile machines
- Image and video monitors
- Wireless terminals
- Monitoring and control equipment
- Bridges, routers, and gateways
- Integrated Services devices, including ISDN terminals and end systems supporting combined voice, video, and data applications

1.4 An international family of standards

The terms LAN and MAN encompass a number of data communications technologies and applications of these technologies. So it is with the IEEE 802® Standards. In order to provide a balance between the proliferation of a very large number of different and incompatible local and metropolitan networks, on the one hand, and the need to accommodate rapidly changing technology and to satisfy certain applications or cost goals, on the other hand, several types of medium access technologies are currently defined in the family of IEEE 802® Standards. In turn, these medium access control (MAC) standards are defined for a variety of physical media. A logical link control (LLC) standard, a secure data exchange standard, and MAC bridging standards are intended to be used in conjunction with the MAC standards. In some ISLAN and MAN standards, provisions are made for optionally conveying isochronous bearer services in support of continuous voice, video, and synchronous data applications. An architecture and protocols for the management of IEEE 802® LANs are also defined.

The IEEE 802® Standards have been developed and applied in the context of an increasingly global data communications industry. This global context is recognized in that most IEEE 802® Standards are progressed to become also international standards, within ISO/IEC JTC 1 (Joint Technical Committee 1, Information Technology, of the International Organization for Standardization and the International Electrotechnical Commission); see Clause 2 and Annex A.

2. References

The following publications contain provisions which, through reference in this text, constitute provisions of this standard. At the time of publication, the editions indicated were valid. All standards are subject to revision, and parties to agreements based on this standard are encouraged to investigate the possibility of applying the most recent editions of the standards listed below.

ISO/IEC 7498-1:1994, Information technology—Open Systems Interconnection—Basic Reference Model: The basic model.¹

ISO/IEC 8802-2:1998 (IEEE Std 802.2-1998[™]), Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 2: Logical Link Control.

ISO/IEC TR 11802-1:1997, Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Technical reports and guidelines—Part 1: The structure and coding of Logical Link Control addresses in Local Area Networks.

ISO/IEC TR 11802-2:1999, Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Technical reports and guidelines—Part 2: Standard Group MAC Addresses.

ISO/IEC 15802-1:1995, Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Common specifications—Part 1: Medium access control (MAC) service definition.

ISO/IEC 15802-3:1998 (IEEE Std 802.1D-1998[™]), Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Common specifications—Part 3: Media Access Control (MAC) Bridges.

NOTE—Annex A lists, for information, the other LAN/MAN and related standards. Unlike those listed above, they do not contain detailed provisions that are used, by reference, by this standard.

3. Definitions, acronyms, and abbreviations

3.1 Definitions

For the purposes of this standard, the following definitions apply. *The Authoritative Dictionary of IEEE Standards Terms*, Seventh Edition [B2],² should be referenced for terms not defined in this clause.

3.1.1 access domain: A set of LAN or MAN stations together with interconnecting data transmission media and related equipment (e.g., connectors, repeaters), in which the LAN or MAN stations use the same MAC protocol to establish the sequence of stations that are in temporary control of the shared transmission media.

3.1.2 bit-reversed representation: The representation of a sequence of octet values in which the values of the individual octets are displayed in order from left to right, with each octet value represented as a two-digit hexadecimal numeral, and with the resulting pairs of hexadecimal digits separated by colons. The order of the hexadecimal digits in each pair, and the mapping between the hexadecimal digits and the bits of the octet value, are derived by reversing the order of the bits in the octet value and interpreting the resulting bit sequence as a binary numeral using the normal mathematical rules for digit significance.

NOTE—The bit-reversed representation is applicable to LAN MAC addresses for use in a Token Ring (IEEE 802.5[®]) or FDDI environment. See Figure 8 for a comparative example of bit-reversed and hexadecimal representation.

¹ISO/IEC publications are available from the ISO Central Secretariat, Case Postale 56, 1 rue de Varembe, CH-1211, Genève 20, Switzerland/Suisse (<http://www.iso.ch/>). ISO/IEC publications are also available in the United States from Global Engineering Documents, 15 Inverness Way East, Englewood, Colorado 80112, USA (<http://global.ihs.com/>). Electronic copies are available in the United States from the American National Standards Institute, 25 West 43rd Street, 4th Floor, New York, NY 10036, USA (<http://www.ansi.org/>).

²The numbers in brackets correspond to those of the bibliography in Annex A.

3.1.3 bridge, MAC bridge: A functional unit that interconnects two or more LANs or MANs that use the same Data Link layer protocols above the MAC sublayer, but can use different MAC protocols.

3.1.4 canonical format: The format of a MAC data frame in which the octets of any MAC addresses conveyed in the MAC user data field have the same bit ordering as in the Hexadecimal Representation.

3.1.5 end station: A device attached to a LAN or MAN, which acts as a source of, and/or destination for, data traffic carried on the LAN or MAN.

3.1.6 Ethernet frame: A MAC data frame structured in accordance with ISO/IEC 8802-3 and containing an Ethernet type value in the LENGTH / TYPE field.

3.1.7 fibre distributed data interface (FDDI) frame: A MAC data frame structured in accordance with ISO/IEC 9314-2.

3.1.8 hexadecimal representation: The representation of a sequence of octet values in which the values of the individual octets are displayed in order from left to right, with each octet value represented as a two-digit hexadecimal numeral, and with the resulting pairs of hexadecimal digits separated by hyphens. The order of the hexadecimal digits in each pair, and the mapping between the hexadecimal digits and the bits of the octet value, are derived by interpreting the bits of the octet value as a binary numeral using the normal mathematical rules for digit significance.

NOTE—See Figure 8 for a comparative example of bit-reversed and hexadecimal representation.

3.1.9 IEEE 802.3® frame, 802.3® frame: A MAC data frame structured in accordance with ISO/IEC 8802-3 and containing a length value in the LENGTH / TYPE field.

3.1.10 IEEE 802.5® frame, 802.5® frame: A MAC data frame structured in accordance with ISO/IEC 8802-5.

3.1.11 IEEE 802.n® frame, 802.n® frame: A MAC data frame structured in accordance with ISO/IEC 8802-n.

NOTES

1—At the time of publication of this standard, relevant specifications, in addition to those cited explicitly in 3.1.9 and 3.1.10, are for n = 4, 6, 9, and 11.

2—ISO/IEC 8802-12 also defines a MAC protocol, but it does not specify its own MAC data frame format; instead, it uses the IEEE 802.3® and IEEE 802.5® frame formats.

3.1.12 IEEE 802® LAN, 802® LAN: A LAN consisting of an access domain using either a MAC protocol specified in one of the IEEE 802.n and ISO/IEC 8802-n Standards or the FDDI MAC protocol.

3.1.13 IEEE 802® MAN, 802® MAN: A MAN consisting of one or more interconnected subnetworks each using a MAC protocol specified in an IEEE 802® or ISO/IEC 8802 MAN Standard.

NOTE—Part of the data communication capability of an IEEE 802® MAN is the provision of a data service equivalent to that provided by an IEEE 802® LAN, over the extended geographical area of the MAN.

3.1.14 interconnection: The provision of data communication paths between LAN or MAN stations.

3.1.15 interworking: The use of interconnected LAN or MAN stations for the exchange of data, by means of protocols operating over the underlying data transmission paths.

3.1.16 LAN: A computer network, located on a user's premises, within a limited geographical area.

3.1.17 MAC control frame: A data structure consisting of fields in accordance with a MAC protocol, for the communication of control information, only, in a LAN or MAN.

NOTE—ISO/IEC 8802-5[®] uses the term “MAC frame” in this sense.

3.1.18 MAC data frame: A data structure consisting of fields in accordance with a MAC protocol, for the communication of user data and control information in a LAN or MAN; one of the fields contains a sequence of octets of user data.

3.1.19 MAC protocol: The protocol that governs access to the transmission medium in a LAN or MAN, to enable the exchange of data between LAN or MAN stations.

3.1.20 MAN: A computer network, extending over a large geographical area such as an urban area and providing integrated communication services such as data, voice, and video.

3.1.21 noncanonical format: The format of a MAC data frame in which the octets of MAC addresses conveyed in the MAC user data field have the same bit ordering as in the Bit-reversed representation.

3.1.22 octet: A sequence of eight bits, the ends of the sequence being identified as the most significant bit (MSB) and the least significant bit (LSB).

NOTE—This identification of the ends of the sequence defines an unambiguous mapping from octet values, via binary numerals, to the integers 0–255, and hence a mapping also from octet values to the expressions of those integers as numerals in hexadecimal notation. See: Hexadecimal Representation.

3.1.23 station: An end station or bridge.

4. Abbreviations and acronyms

CMIP	common management information protocol (ISO/IEC 9596-1)
CSMA/CD	carrier sense multiple access with collision detection (ISO/IEC 8802-3)
DQDB	distributed queue dual bus
FDDI	fibre distributed data interface (ISO/IEC 9314)
IM	implementation model
I/G	individual/group
ISDN	integrated services digital network
ISLAN	integrated services LAN (ISO/IEC 8802-9)
LAN	local area network
LLC	logical link control (ISO/IEC 8802-2)
LSAP	link service access point (ISO/IEC 8802-2)
LSB	least significant bit
MAC	medium access control, media access control ³
MAN	metropolitan area network

³Both forms are used, with the same meaning. This standard uses “medium.”

MIB	management information base
MOCS	managed object conformance statement
MSAP	MAC service access point
MSB	most significant bit
PDU	protocol data unit
PhSAP	physical service access point
PICS	protocol implementation conformance statement
RM	reference model
OSI	open systems interconnection (ISO/IEC 7498-1)
OUI	organizationally unique identifier
SNAP	subnetwork access protocol
SNMP	simple network management protocol (RFC 1157 ⁴)
U/L	universally or locally administered
WAN	wide area network

5. Compliance

NOTE—IEEE policy with respect to claims of compliance, conformance, or compatibility with IEEE standards can be found in the Trademark Policy statement in the front matter. This clause will be deleted in the next revision of this standard.

6. Reference and implementation models

6.1 Introduction

This clause defines the IEEE 802[®] LAN and MAN RM (LAN&MAN/RM) and implementation model (LAN&MAN/IM). The intent of presenting these models is as follows:

- a) To provide an overview of the standard
- b) To serve as a guide to reading other IEEE 802[®] Standards

The IEEE 802[®] LAN&MAN/RM is patterned after the Open Systems Interconnection (OSI) Basic Reference Model (OSI/RM), ISO/IEC 7498-1. It is assumed that the reader has some familiarity with the OSI/RM and its terminology. The IEEE 802[®] Standards encompass the functionality of the lowest two layers of the OSI/RM (i.e., Physical layer and Data Link layer) and the higher layers as they relate to LAN management. The LAN&MAN/RM is similar to the OSI/RM in terms of its layers and the placement of its service boundaries.

For the mandatory packet services supported by all LANs and MANs, the Data Link layer is structured as two sublayers, with the LLC sublayer operating over a MAC sublayer. In addition, some IEEE 802[®] LAN technologies provide direct support by the MAC sublayer for an alternative Ethernet sublayer operating at the same place in the architecture as does LLC; for the other IEEE 802[®] LAN technologies, the equivalent

⁴See A.1 for information on obtaining RFCs.

functionality is provided by encapsulation of the Ethernet sublayer information within LLC Protocol Data Units (PDUs), using the Subnetwork Access Protocol specified in Clause 10 of this standard.

Optionally, some Integrated Service LANs and MANs may also support ITU-T compatible isochronous bearer services at the Physical layer.

The OSI/RM is referred to by the IEEE 802[®] Standards because of the following:

- The OSI/RM provides a common vehicle for understanding and communicating the various components and interrelationships of the standards.
- The OSI/RM helps define terms.
- The OSI/RM provides a convenient framework to aid in the development and enhancement of the standards.
- The use of the OSI/RM facilitates a higher degree of interoperability than might otherwise be possible.

Figure 1 shows the architectural view of LAN&MAN/RM and its relation to the OSI/RM.

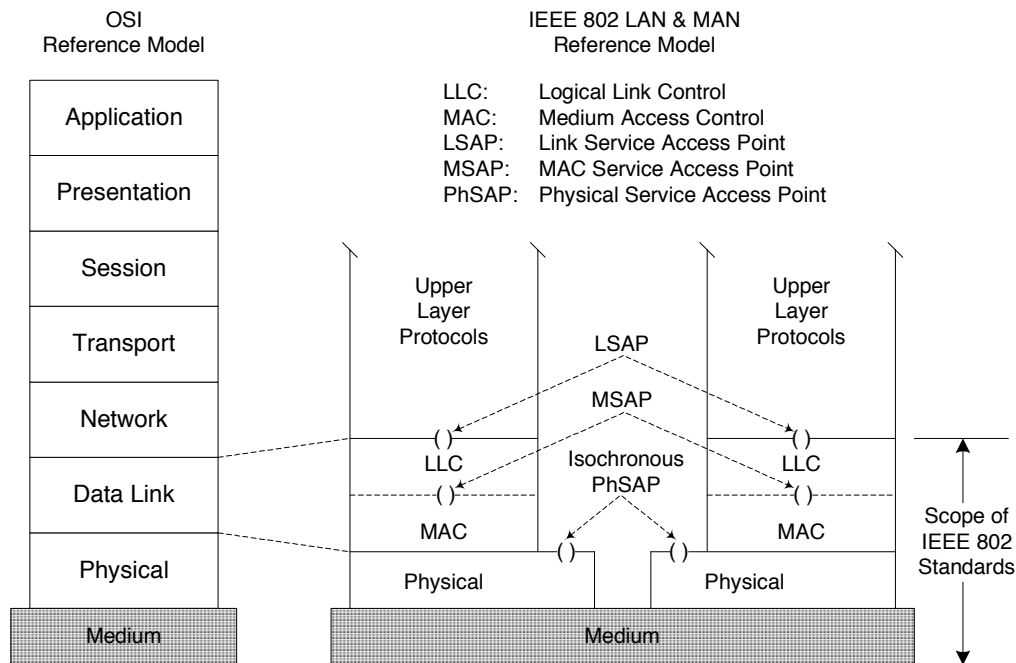


Figure 1—IEEE 802[®] RM for end stations (LAN&MAN/RM)

The LAN&MAN/IM is more specific than the LAN&MAN/RM, allowing differentiation between implementation approaches (e.g., of carrier sense multiple access with collision detection [CSMA/CD], and token passing). Figure 2 shows two implementation views of LAN&MAN/IMs and their relation to the LAN&MAN/RM.

Both Figure 1 and Figure 2 illustrate the application of the models in LAN end stations. A variation of the model applies within bridges (see 6.3.2). Also, Figure 1 and Figure 2 illustrate only the basic transfer of user data between end stations.

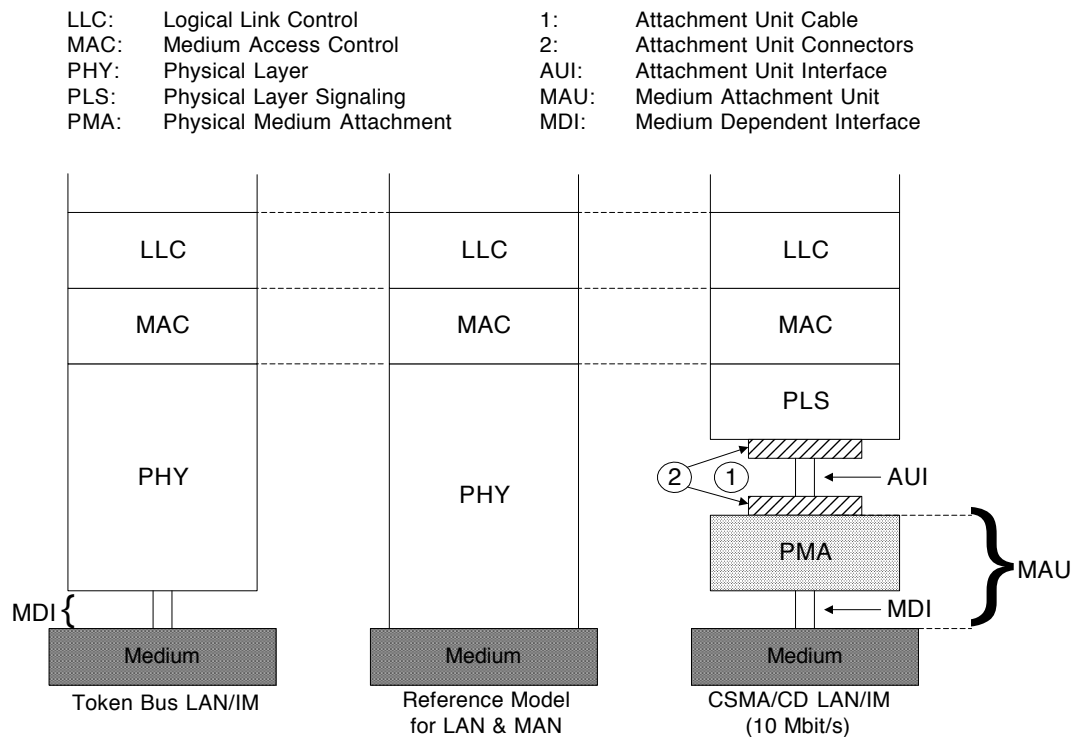


Figure 2—IEEE 802® RM and two examples of end-station implementation model

Considerations of management and security in LANs and MANs are also covered by IEEE 802® standards; these optional features lead to an elaboration of the RM (Figure 3). LAN/MAN management provides a Data Link layer management protocol for exchange of management information between LAN stations; managed objects are defined for all LAN/MAN protocol standards. The Secure Data Exchange (SDE) entity forms part of the LLC sublayer and provides a secure connectionless service immediately above the MAC sublayer.

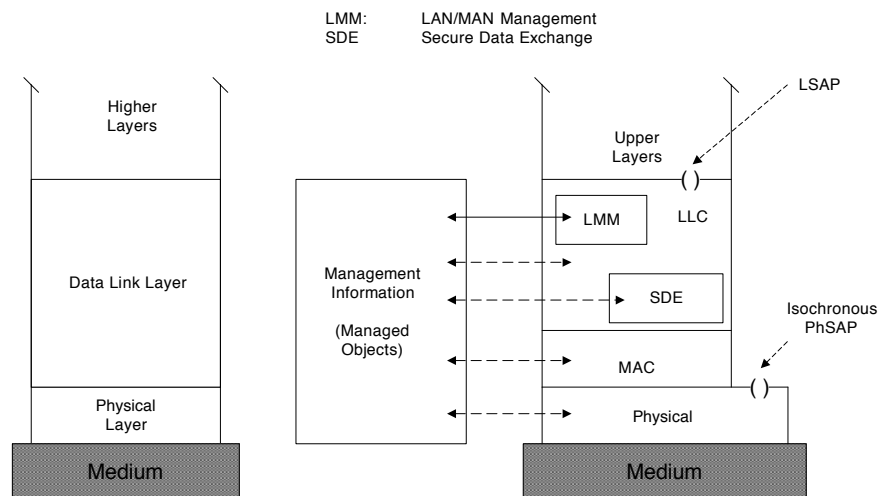


Figure 3—IEEE 802® RM with end-station management and security

6.2 RM description for end stations

The LAN&MAN/RM maps to the OSI/RM as shown in Figure 1. The applicable part of the OSI/RM consists of the lowest two layers: the Data Link layer and the Physical layer. These map onto the same two layers in the IEEE LAN&MAN/RM. The MAC sublayer of the LAN&MAN/RM exists between the Physical layer and the LLC sublayer to provide a common service for the LLC sublayer (certain MAC types provide additional MAC service features that can be used by LLC, in addition to the common core features). Service access points (SAPs) for addressing endpoints are shown.

6.2.1 Service access points (SAPs)

Multiple-link service access points (LSAPs) provide interface ports to support multiple higher layer users above the LLC sublayer.

The MAC sublayer provides a single MAC service access point (MSAP) as an interface port to the LLC sublayer in an end station. In general, the MSAP is identified (for transmission and reception) by a single individual MAC address and (for reception) by the LAN-wide broadcast MAC address; it can also be identified (for reception) by one or more group MAC addresses. Clause 9 provides details of how these MAC addresses are constructed and used; see also ISO/IEC 15802-1.

A user of LLC is identified by, at a minimum, the logical concatenation of the MAC address field and the LLC address field in a frame. See ISO/IEC 8802-2 and ISO/IEC TR 11802-1 for a description of LLC addresses.

The Physical layer provides an interface port to a single MAC station, and in the case of ISLANs and MANs, it may optionally offer isochronous bearer services at multiple Physical service access points (PhSAPs). See 11.3 for a more detailed description of ISLAN and MAN PhSAP addressing requirements.

6.2.2 LLC sublayer

The LLC sublayer standard, ISO/IEC 8802-2, describes three types of operation for data communication between service access points: unacknowledged connectionless-mode (type 1), connection-mode (type 2), and acknowledged connectionless-mode (type 3).

With type 1 operation, information frames are exchanged between LLC entities without the need for the prior establishment of a logical link between peers. The LLC sublayer does not provide any acknowledgments for these LLC frames, nor does it provide any flow control or error recovery procedures.

LLC type 1 also provides a TEST function and an Exchange Identification (XID) function. The capability to act as responder for each of these functions is mandatory: This allows a station that chooses to support initiation of these functions to check the functioning of the communication path between itself and any other station, to discover the existence of other stations, and to find out the LLC capabilities of other stations.

With type 2 operation, a logical link is established between pairs of LLC entities prior to any exchange of information frames. In the data transfer phase of operation, information frames are transmitted and delivered in sequence. Error recovery and flow control are provided, within the LLC sublayer.

With type 3 operation, information frames are exchanged between LLC entities without the need for the prior establishment of a logical link between peers. However, the frames are acknowledged to allow error recovery and proper ordering. Further, type 3 operation allows one station to poll another for data.

NOTE—ISO/IEC 8802-2 defines four classes of LLC, each of which groups together support for a different combination of LLC types. All classes include mandatory support of type 1.

The Secure Data Exchange (SDE) entity forms part of the LLC sublayer and provides a secure connectionless service immediately above the MAC sublayer. The operation of the SDE entity is described in IEEE Std 802.10®.

6.2.3 MAC sublayer

The MAC sublayer performs the functions necessary to provide packet-based, connectionless-mode (datagram style) data transfer between stations in support of the LLC sublayer or of the Ethernet sublayer (see 6.1) for LANs that support it. The term *MAC frame*, or simply *frame*, is used to describe the packets transferred within the MAC sublayer. In some MAC types (e.g., Token Ring), some MAC frames are used in support of the MAC sublayer functionality itself, rather than for transfer of LLC data.

The principal functions of the MAC sublayer comprise the following:

- Frame delimiting and recognition
- Addressing of destination stations (both as individual stations and as groups of stations)
- Conveyance of source-station addressing information
- Transparent data transfer of LLC PDUs, or of equivalent information in the Ethernet sublayer
- Protection against errors, generally by means of generating and checking frame check sequences
- Control of access to the physical transmission medium

Other functions of the MAC sublayer—applicable particularly when the supporting implementation includes interconnection devices such as hubs or bridges—include flow control between an end station and an interconnection device (see 6.3), and filtering of frames according to their destination addresses to reduce the extent of propagation of frames in parts of a LAN or MAN that do not contain communication paths leading to the intended destination end station(s).

The functions listed are those of the MAC sublayer as a whole. Responsibility for performing them is distributed across the transmitting and receiving end stations, and any interconnection devices such as bridges. Devices with different roles therefore can behave differently in support of a given function. For example, transmission of a MAC frame by a bridge is very similar to transmission by an end station, but not identical. Principally, the handling of source-station addressing is different.

The various MAC specifications all specify MAC frame formats in terms of a serial transmission model for the service provided by the supporting Physical layer. This model supports concepts such as “first bit (e.g., of a particular octet) to be transmitted,” and a strict order of octet transmission, in a uniform manner. However, the ways in which the model has been applied in different MAC specifications are not completely uniform with respect to bit-ordering within octets (see Clause 9, and particularly 9.5, for examples and explanation).

NOTE—The serial transmission model does not preclude current or future MAC specifications from using partly or wholly octet-oriented specifications of frame formats or of the interface to the Physical layer.

6.2.4 Physical layer

The Physical layer provides the capability of transmitting and receiving bits between Physical layer entities. A pair of Physical layer entities identifies the peer-to-peer unit exchange of bits between two MAC users, and in the case of ISLANs and MANs, it may optionally support the exchange of bits with end-to-end timing preserved between isochronous-service PhSAPs.

The Physical layer provides the capability of transmitting and receiving modulated signals assigned to specific frequency channels, in the case of broadband or wireless, or to a single-channel band, in the case of baseband.

Note that, whereas the service offered to the MAC sublayer is expressed as the transfer of bits (in sequences representing MAC frames), the actual symbols that are encoded for transmission do not always represent individual bits. Particularly at speeds of 100 Mbit/s and above, or for wireless transmission, the Physical layer can map blocks of several bits (e.g., 4, 5, or 8 bits) to different multi-element symbols. In some Physical-layer encodings, these symbols are subject to further transformation before transmission, and in some cases, the transmission is spread over multiple physical data paths. The actual transmission on physical media can therefore be far removed from the simple bit-serial representation of a MAC frame (as was specified, for example, in the original ISO/IEC 8802-3 and ISO/IEC 8802-5 Physical layers).

6.2.5 Layer and sublayer management

The LLC sublayer, MAC sublayer, and Physical layer standards also include a management component that specifies managed objects and aspects of the protocol machine that provides the management view of these resources. See Clause 8 for further information.

6.3 Interconnection and interworking

In some cases, the end systems on a LAN or MAN have no need to communicate with end systems on other networks (other LANs, WANs, etc.). However, this is not expected to be the norm; there are many cases in which end stations on a LAN or MAN will need to communicate with end systems on other networks and so devices that interconnect the LAN or MAN with other kinds of networks are required. In addition, several standard methods have been developed that permit a variety of interconnection devices to operate transparently to end stations on a LAN or MAN in order to extend the LAN/MAN capabilities available to end stations, particularly in terms of the geographical extent and/or total number of end stations that can be supported.

Standard methods of interworking fall into the following three general categories, depending on the layer at which the corresponding interconnection devices operate:

- Physical-layer interconnection, using devices usually termed *repeaters* or *hubs* (6.3.1)
- MAC-sublayer interconnection, using devices termed *bridges* (6.3.2)
- Network-layer interconnection, using devices usually termed *routers* (6.3.3)

See also Clause 11 of this standard, for an outline of the optional methods by which ISLANs and MANs may support isochronous interworking with WANs and with remote ISLANs and MANs.

6.3.1 Physical-layer interconnection: repeaters and hubs

The original IEEE 802® LAN specifications were for end stations attached to a shared communication medium. This basic LAN configuration is referred to as a single *access domain*; the domain consists of the set of LAN stations such that at most one can be transmitting at a given time, with all other stations acting as (potential) receivers.

A repeater is a device used to interconnect segments of the physical communications media, for example, to extend the range of a LAN when the physical specifications of the technology would otherwise be exceeded, while providing a single access domain for the attached LAN stations. Repeaters used in support of multiple end stations attached by star-wired network topologies are frequently referred to as hubs.

6.3.2 MAC-sublayer interconnection: bridges

6.3.2.1 Bridges and bridged LANs

Bridges (see 3.1.3) are devices that interconnect multiple access domains. ISO/IEC 15802-3 provides the basic specification for bridge interworking among IEEE 802® networks. A *bridged LAN* (see 3.1 of ISO/IEC 15802-3) consists of one or more bridges together with the complete set of access domains that they interconnect. A bridged LAN provides end stations belonging to any of its access domains with the appearance of a LAN that contains the whole set of attached end stations.

A bridged LAN can provide for the following:

- Communication between stations attached to LANs of different MAC types
- An increase in the total throughput of a LAN
- An increase in the physical extent of, or number of permissible attachments to, a LAN
- Partitioning of the physical LAN for administrative or maintenance reasons

The term *switch* is often used to refer to some classes of bridge. However, there is no consistent meaning applied to the distinction between the terms bridge and switch, and ISO/IEC 15802-3 does not make any such distinction. Hence, this standard only uses the term bridge.

6.3.2.2 Relaying and filtering by bridges

A bridge processes protocols in the MAC sublayer and is functionally transparent to LLC and higher layer protocols. MAC frames are forwarded between access domains, or filtered (i.e., not forwarded to certain access domains), on the basis primarily of MAC addressing information. Figure 4 shows the position of the bridging functions within the MAC sublayer; note particularly that the relaying and filtering functions are considered to belong entirely within the MAC sublayer.

Filtering by bridges tends to confine traffic to only those parts of the bridged LAN that lie between transmitting end stations and the intended receivers. This permits a bridged LAN to support several transmitting end stations at any given time (up to the total number of access domains present).




Figure 4—Internal organization of the MAC sublayer with bridging

6.3.2.3 The Spanning Tree Protocol

A key aspect of ISO/IEC 15802-3 is its specification of the Spanning Tree Protocol, which is used by bridges to configure their interconnections in order to prevent looping data paths in the bridged LAN. In the event that the basic interconnection topology of bridges and LANs contains multiple possible paths between certain points, use of the Spanning Tree Protocol blocks some paths in order to produce a simply connected

active topology for the flow of MAC user traffic between end stations. For each point of attachment of a bridge to a LAN, the Spanning Tree Protocol selects whether or not MAC user traffic is to be received and transmitted by the bridge at that point of attachment.

The Spanning Tree Protocol adapts to changes in the configuration of the bridged LAN, maintaining connectivity while avoiding data loops. Some configuration changes can cause temporary interruptions of connectivity between parts of the bridged LAN, typically lasting for a few tens of seconds at most.

6.3.2.4 Transparent bridging and source routing

ISO/IEC 15802-3 specifies *transparent bridging* operation, so called because the MAC bridging function does not require the MAC user frames transmitted and received to carry any additional information relating to the operation of the bridging functions; end-station operation is unchanged by the presence of bridges.

Also specified, in an annex to ISO/IEC 15802-3, is the operation of a Source Routing Transparent (SRT) bridge. This adds the ability for a *source-routing* function in a bridge to use explicit MAC-sublayer routing information contained in MAC user frames; this source-routing information is inserted by the transmitting end station. Source routing is specified only for ISO/IEC 8802-5 Token Ring and ISO/IEC 9314 FDDI LANs, and for IEEE 802.12® Demand Priority LANs when operating with 8802-5 format frames.

6.3.2.5 Remote MAC bridging

ISO/IEC 15802-5 extends the specification in ISO/IEC 15802-3 to cover remote MAC bridging, where non-IEEE 802® communications technologies are used to interconnect bridges, for example, to allow a bridged LAN to include WAN links to provide greatly extended geographical range. Figure 5 shows the corresponding organization of the MAC sublayer.




Figure 5—Internal organization of MAC sublayer with remote bridges

6.3.2.6 Bridging example

Some bridges are used to interconnect access domains each containing a very small number of end stations (often, a single end station). Others interconnect multiple access domains containing principally other bridges, thus, forming a backbone for the bridged LAN. (Hybrid configurations, with characteristics of both kinds of bridge, are of course possible.) Bridged LAN configurations involving these kinds of interconnection have become widespread as the IEEE 802® LAN technologies have developed. As noted in 1.2, they allow the construction of networks with much larger numbers of end stations and much higher aggregate throughput than was previously achievable.

Figure 6 illustrates the kind of bridged LAN that can be configured with bridge-style interconnection. The bridges A and B, and the CSMA/CD LAN configurations to which they attach, are typical of the older style of bridged LAN in which a bridge interconnects a small number of access domains each containing many

end stations, as is similar with K, L, and M and their Token Ring LANs. On the other hand, the bridges S, T, and U function as bridges in a high-speed backbone that combines FDDI and 100-Mbit/s IEEE 802.3® LANs. S is a backbone bridge, handling a number of high-speed LAN attachments. T and U are bridges that support multiple end stations, with connection to the backbone. B and K also provide access to the backbone. The end station shown connected to S by a point-to-point link could be a server system, as might the end stations attached to the FDDI LAN.




Figure 6—A bridged LAN

6.3.3 Network-layer interconnection: routers

The third category of interconnection uses Network-layer interconnection devices, generally known as routers, that operate as LAN/MAN end stations. These process Network layer protocols that operate directly above the LLC sublayer or equivalent, with forwarding decisions based on Network layer addresses. Details of this kind of interconnection lie outside of the scope of IEEE 802® LAN/MAN Standards, but the various standard and proprietary Network-layer protocols involved represent a very substantial part of the user traffic on many real IEEE 802® networks. In particular, IEEE 802® LANs and MANs are often interconnected by routers for the Internet Protocol (IP) and its related routing and management protocols, either directly to other LANs or by means of WAN links.

7. General requirements for an 802® LAN or MAN

7.1 Services supported

With the descriptions in Clause 6 as a basis, an IEEE 802® LAN or MAN can be characterized as a communication resource that provides sufficient capabilities to support the MAC service defined in ISO/IEC 15802-1, between two or more MSAPs. In particular, this requires the ability to convey LLC data from one MSAP to n other MSAPs, where n can be any number from 1 to all of the other MSAPs on the network. An IEEE 802® LAN is required, at a minimum, to support both LLC Type 1 and the MAC Internal Sublayer Service defined in ISO/IEC 15802-3. In addition, an ISLAN or MAN may optionally support isochronous bearer services compatible with ISDN services as defined in the ITU-T I-series Recommendations.

7.2 Size and extent

The original IEEE 802® LAN and MAN technologies were designed to be capable of supporting access domains containing at least 200 end stations and with geographical extent of at least 2 km for LANs (using Physical-layer repeaters if necessary) and 50 km for MANs. Subsequent developments in IEEE 802® LAN technology and performance have been accompanied by a reduction in the size and extent required in individual access domains, recognizing that these can readily and cost-effectively be interconnected in bridged LANs that are capable of offering at least the original minimum size and extent, with increased overall bandwidth and performance. Size and extent requirements for future IEEE 802® LAN technologies are, similarly, expected to be determined by application needs and opportunities.

7.3 Error rates

Error performance of IEEE 802® LANs and MANs is required to be such as follows:

- a) For wired or optical fiber physical media: Within a single access domain, the probability that a transmitted MAC frame (excluding any preamble) is not reported correctly at the Physical Service interface of an intended receiving peer MAC entity, due only to operation of the Physical layer, shall be less than 8×10^{-8} per octet of MAC frame length.
- b) For wireless physical media: Within a single access domain, the probability that a MAC Service Data Unit (MSDU) is not delivered correctly at an MSAP of an intended receiving MAC service user, due to the operation of the Physical layer and the MAC protocol, shall be less than 8×10^{-8} per octet of MSDU length.

NOTE—The performance measure stated in (a) defines a highly desirable characteristic of LAN performance, as it has a bearing on other aspects of the delivered service, such as frame loss and transmission delays caused by the need to retransmit. However, this measure is not realistic for all physical media; for example, wireless media may be unable to meet this level of physical layer performance due to the inherent transmission characteristics of the medium. In such cases, the operation of the MAC protocol must employ additional mechanisms, for example, error detection and correction mechanisms, in order to enable the MAC service provider to meet the performance levels implied by this condition in the service offered at the MAC service boundary.

- c) The probability that an MSDU delivered at an MSAP contains an undetected error, due to operation of the MAC service provider, shall be less than 5×10^{-14} per octet of MSDU length.

NOTE—For example, (a) the worst-case probability of losing a maximum-length IEEE 802.3® frame (1518 octets) through physical-layer damage is to be less than 1.21×10^{-4} , or approximately 1 in 8250; (c) the worst-case probability that a similar frame, which contains an MSDU of 1500 octets, is delivered with an undetected error is to be less than 7.5×10^{-11} , or approximately 1 in 13 300 000 000.

7.4 Service availability

Insertion of a device into, or removal of a device from, a LAN or MAN shall cause at most a transient loss of availability of the access domain(s) to which the device attaches, lasting not more than 1 s. Failure of a device, including loss of power, shall not cause more than a transient fault for the access domain(s) to which it attaches, with duration of order 1 s.

NOTE—In a bridged LAN, reconfiguration of the topology in response to logical insertion or removal of a bridge, or to changes in a bridge's configuration parameters, can cause loss of communication between some access domains for longer periods, typically a few tens of seconds; ISO/IEC 15802-3 contains the full specification.

7.5 Safety, and lightning and galvanic protection

Equipment implementing IEEE 802® LAN and MAN standards is typically subject to guidance and requirements relating to safety and to protection of the equipment and its users from lightning and galvanic effects. Such guidance and requirements are outside of the scope of IEEE 802® standardization; they are

typically specified by other organizations with different legal, geographical, and industrial scope. However, the general underlying concerns can have an influence on the Physical layer aspects of IEEE 802® LAN and MAN standards.

7.6 Regulatory requirements

Equipment implementing IEEE 802® LAN and MAN standards may be subject to regulations imposed within particular geographical and political domains. For example, the deployment of equipment implementing IEEE 802® wireless LAN standards may be subject to local regulations that pertain to the use of radio-frequency transmission. Such regulations are outside of the scope of IEEE 802® standardization; they are typically specified by other organizations with different legal, geographical, and industrial scope. However, the general underlying concerns can have an influence on the Physical layer aspects of IEEE 802® LAN and MAN standards.

8. LAN/MAN management

The provision of an adequate means of remote management is an important factor in the design of today's LAN equipment. Such management mechanisms fall into two broad categories: those that provide general-purpose management capability, allowing control and monitoring for a wide variety of purposes, and those that provide specific capabilities aimed at a particular aspect of management. These aspects of management are discussed in 8.1 and 8.2, respectively.

8.1 General-purpose LAN/MAN management

This subclause introduces the functions of management to assist in the identification of the requirements placed on IEEE 802® LAN/MAN equipment for support of management facilities, and it identifies general-purpose management standards that may be used as the basis of developing management specifications for such equipment.

8.1.1 Management functions

Management functions relate to users' needs for facilities that support the planning, organization, supervision, control, protection and security of communications resources, and account for their use. These facilities may be categorized as supporting the functional areas of configuration, fault, performance, security, and accounting management. These can be summarized as follows.

- Configuration management provides for the identification of communications resources, initialization, reset and close-down, the supply of operational parameters, and the establishment and discovery of the relationships between resources.
- Fault management provides for fault prevention, detection, diagnosis, and correction.
- Performance management provides for evaluation of the behavior of communications resources and of the effectiveness of communication activities.
- Security management provides for the protection of resources.
- Accounting management provides for the identification and distribution of costs and the setting of charges.
- Management facilities in LAN/MAN equipment will address some or all of these areas, as appropriate to the needs of that equipment and the environment in which it is to be operated.

8.1.2 Management architecture

The management facilities defined in IEEE 802[®] LAN and MAN standards are based on the concept of managed objects, which model the semantics of management operations. Operations on a managed object supply information concerning, or facilitate control over, the process or entity associated with that object.

Operations on a managed object can be initiated by mechanisms local to the equipment being managed (e.g., via a control panel built into the equipment), or can be initiated from a remote management system by means of a general-purpose management protocol carried using the data services provided by the LAN or MAN to which the equipment being managed is connected.

There are two general-purpose management protocols of relevance to the management of LAN/MAN equipment, as follows:

- The Simple Network Management Protocol (SNMP), as described in RFC 1157
- The OSI common management information protocol (CMIP), as described in ISO/IEC 9595 and ISO/IEC 9596 and related standards

NOTE—In addition to operation of CMIP over a full OSI protocol stack, two standards define the use of CMIP over simpler protocol support. ISO/IEC 15802-2 (IEEE Std 802.1B-1995[™]) defines a means of using CMIP over a simple Data Link layer protocol stack in LANs; RFC 1095 describes the use of CMIP over a TCP/IP protocol stack.

Of the two protocols, SNMP is the more significant in terms of its wide application across the spectrum of LAN/MAN products in today's marketplace; however, in some markets, and where it is desirable to integrate LAN management with management of wide-area networking and telephony equipment, use of the OSI management protocols may be important.

8.1.3 Managed object definitions

In order for an IEEE 802[®] standard to specify management facilities, it is necessary for them to define managed objects that model the operations that can be performed on the communications resources specified in the standard. There are essentially two components to a managed object definition, as follows:

- 1) A definition of the functionality provided by the managed object, and the relationship between this functionality and the resource to which it relates
- 2) A definition of the syntax that is used to convey management operations, and their arguments and results, in a management protocol

The functionality of a managed object can be described in a manner that is independent of the protocol that will be used; this abstract definition can then be used in conjunction with a definition of the syntactic elements required in order to produce a complete definition of the object for use with specific management protocols.

Each management protocol has its notation for defining managed objects, as follows:

- a) SNMP has standards for the structure of management information known as SMIV1 (RFC 1155, RFC 1212 and RFC 1215) and SMIV2 (RFC 1902, RFC 1903 and RFC 1904), which provides ASN.1-based macros for defining managed objects.
- b) CMIP has a standard language, known as GDMO (ISO/IEC 10165-4), which is used along with ASN.1 to define both the syntactic and semantic aspects of managed objects.

The choice of notational tools for defining managed objects will depend on which of the available management protocols the standard will support.

NOTES

1—IEEE Std 802.1F[®] provides additional guidance for use of GDMO in LAN/MAN standards.

2—Some IEEE 802[®] standards have used GDMO as the notation for their managed object definitions, with SNMP management information base (MIB) definitions being developed subsequently within the IETF, using automatic tools for translating GDMO definitions into equivalent SNMP definitions.

8.2 Special-purpose LAN/MAN management standards

Special-purpose protocols relating to the management functionality of IEEE 802[®] stations can be developed where the use of a general-purpose management protocol is inappropriate. An example of a special-purpose management protocol is ISO/IEC 15802-4, which defines the services and protocols for remote station loading in a LAN/MAN environment. This protocol permits the simultaneous loading of multiple stations by use of the group-addressing capability in IEEE 802[®] technologies.

9. Universal addresses and protocol identifiers

The IEEE makes it possible for organizations to employ unique individual LAN MAC addresses, group addresses, and protocol identifiers. It does so by assigning Organizationally Unique Identifiers (OUIs), which are three octets (24 bits) in length. Because the assignment of the OUI in effect reserves a block of each derivative identifier (i.e., blocks of individual LAN MAC addresses, group addresses, and protocol identifiers), the address space of the OUI is chosen to be large. Although the OUIs are 24 bits in length, their true address space is 22 bits. The LSB of the first octet can be set to 1 or 0 depending on the application. The next-to-LSB of the first octet is 0, for all assignments. The remaining 22 bits, which shall not be changed by the assignee, result in 2²² (approximately 4 million) identifiers; see Figure 7.




Figure 7—OUI

The universal administration of LAN MAC addresses began with the Xerox Corporation administering Block Identifiers (Block IDs) for Ethernet addresses. Block IDs were assigned by the Ethernet Administration Office and were 24 bits in length (three octets). An organization developed addresses by assigning the remaining 24 bits. For example, the address as represented by the six octets P-Q-R-S-T-U comprises the Block ID, P-Q-R, and the locally assigned octets S-T-U.

The IEEE, because of the work in Project 802[®] on standardizing LAN technologies, has assumed the responsibility of defining and carrying out procedures for the universal administration of addresses for IEEE and ISO/IEC LANs (e.g., CSMA/CD, Token Bus, Token Ring, and FDDI). In carrying out the procedures, the IEEE acts as the Registration Authority for OUIs.⁵ The responsibility for defining the procedures is discharged by the IEEE Registration Authority Committee (RAC), which is chartered by the IEEE Standards Association Board of Governors.

⁵ Interested applicants should contact the IEEE Registration Authority, Institute of Electrical and Electronic Engineers Inc., 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331, USA.

The IEEE honors the Block ID assignments made by the predecessor administration office where those assignments fall—as the great majority of them do—within the space administered by the IEEE. The Block ID is referred to as the OUI by the IEEE.

9.1 OUI

OUIs allow a general means of assuring unique identifiers for a number of purposes. Currently, the IEEE assigns OUIs to be used for generating LAN MAC addresses and protocol identifiers. Assuming correct administration by the IEEE Registration Authority and the assignee, the LAN MAC addresses and protocol identifiers will be universally unique.

OUIs are assigned as three-octet values. Both values (0, 1) are assigned to the LSB of the first octet. The next-to-LSB of the first octet is set to 0; this bit of the OUI being set to 0 indicates that the assignment is universal. Three-octet values occupying the same fields as OUIs can occupy, but with the next-to-LSB of the first octet set to 1, are locally assigned and have no relationship to the IEEE-assigned values (as described herein).

The standard representation of the OUI is as a string of three octets, using the hexadecimal representation (see 3.1.8).

9.2 48-bit universal LAN MAC addresses

9.2.1 Concept

The concept of universal addressing is based on the idea that all potential members of a network need to have a unique identifier (if they are going to coexist in the network). The advantage of a universal address is that a station with such an address can be attached to any LAN in the world with an assurance that the address is unique.

A 48-bit universal address consists of two parts. The first 24 bits correspond to the OUI as assigned by the IEEE, except that the assignee may set the LSB of the first octet to 1 for group addresses or set it to 0 for individual addresses. The second part, comprising the remaining 24 bits, is administered by the assignee. In the 48-bit LAN MAC address, an example of which is shown in Figure 8, the OUI is contained in octets 0, 1, 2, and the value assigned by the assignee is contained in octets 3, 4, 5. This address, including its OUI, is used throughout this document as the basis for examples of LAN MAC addresses and protocol identifiers.

NOTE—The requirement for the use of 64-bit addresses in new applications is under consideration by the IEEE Registration Authority (RAC).

The standard representation of a 48-bit LAN MAC address is as a string of six octets, using the hexadecimal representation (3.1.8). In certain contexts associated with use of IEEE 802.5[®] frame formats, LAN MAC addresses may be represented using the alternative bit-reversed representation (3.1.2). See 9.5 for further specification relating to use of the bit-reversed representation.

NOTE—The upper, bit-stream representation of the universal address in Figure 8 shows the LSB of each octet first; this corresponds to the data-communications convention for representing bit-serial transmission in left-to-right order, applied to the model for transmission of LAN MAC address fields (see 6.2.3). See also 9.5 for further discussion of bit-ordering issues. The lower, octet-sequence representation shows the bits within each octet in the usual order for binary numerals; the order of octet transmission is from the top downward.

Hexadecimal representation: AC-DE-48-00-00-80
 Bit-reversed representation: 35:7B:12:00:00:01




Figure 8—Universal address

The Individual/Group (I/G) address bit (LSB of octet 0) is used to identify the destination address as an individual address or a group address. If the I/G address bit is 0, it indicates that the address field contains an individual address. If this bit is 1, the address field contains a group address that identifies one or more (or all) stations connected to the LAN. The all-stations broadcast address is a special, predefined group address of all 1's.

The Universally or Locally administered (U/L) address bit is the bit of octet 0 adjacent to the I/G address bit. This bit indicates whether the address has been assigned by a local or universal administrator. Universally administered addresses have this bit set to 0. If this bit is set to 1, the entire address (i.e., 48 bits) has been locally administered.

9.2.2 Assignment by organizations

Varying the last 24 bits in the block of MAC addresses for a given OUI allows the OUI assignee approximately 16 million unique individual addresses and 16 million unique group addresses that no other organization may assign (i.e., universally unique). The IEEE intends not to assign additional OUIs to any organization unless the organization has exhausted this address block. Therefore, it is important for the IEEE to maintain a single point of contact with each assignee to avoid complicating the assignment process. It is important to note that in no way should these addresses be used for purposes that would lead to skipping large numbers of them (for example, as product identifiers for the purpose of aiding company inventory procedures). The IEEE asks that organizations not misuse the assignments of the last 24 bits and thereby unnecessarily exhaust the block. There are sufficient identifiers to satisfy most needs for a long time, even in volume production; however, no address space is infinite.

The method that an assignee uses to ensure that no two of its devices carry the same address will, of course, depend on the assignment or manufacturing process, the nature of the organization, and the organization's philosophy. However, the users of networks worldwide expect to have unique addresses. The ultimate responsibility for assuring that user expectations and requirements are met, therefore, lies with the organization offering such devices.

The recommended approach is for each device associated with a distinct point of attachment to a LAN to have its own unique MAC address. Typically, therefore, a LAN adapter card (or, e.g., an equivalent chip or set of chips on a motherboard) should have one unique MAC address for each LAN attachment that it can support at a given time.

9.3 Protocol identifiers

The protocol identifier is five octets (40 bits) in length and follows the LLC header in a frame. The first three octets of the protocol identifier consist of the OUI in exactly the same fashion as in 48-bit LAN MAC addresses. The remaining two octets (16 bits) are administered by the assignee. In the protocol identifier, an example of which is shown in Figure 9, the OUI is contained in octets 0, 1, 2 with octets 3, 4 being assigned by the assignee of the OUI.

							X bit	M bit
Octet 0	1	0	1	0	1	1	0	0
Octet 1	1	1	0	1	1	1	1	0
Octet 2	0	1	0	0	1	0	0	0
Octet 3	0	0	0	0	0	0	0	0
Octet 4	1	0	0	0	0	0	0	0

Copyright © 2002 IEEE. All rights reserved.

The standard representation of a protocol identifier is as a string of five octets, using the hexadecimal representation (3.1.8).

The LSB of the first octet of a protocol identifier is referred to as the M bit. All identifiers derived from OUIs assigned by the IEEE shall have the M bit set to 0. Values with the M bit set to 1 are reserved.

Protocol identifiers may be assigned universally or locally. The X bit of a protocol identifier is the bit of the first octet adjacent to the M bit. All identifiers derived from OUIs assigned by the IEEE will have the X bit set to 0 and are universally assigned. Values with the X bit set to 1 are locally assigned and have no relationship to the IEEE-assigned values. They may be used, but there is no assurance of uniqueness.

9.4 Standard group MAC addresses

The previous subclauses described the assignment of individual and group MAC addresses, and protocol identifiers for public or private use by private organizations. There is also a need for standard group MAC addresses to be used with standard protocols. The administration of these addresses, including the procedure for application and a list of currently assigned values, is defined in ISO/IEC TR 11802-2. These standard group MAC addresses come from a block of universally administered LAN MAC addresses derived from an OUI that has been assigned by the IEEE for this purpose.

ISO/IEC 8802-5 also defines *functional addresses*, for use in Token Ring LAN environments to identify well-known functional entities. These addresses are a subset of the locally administered group MAC addresses, identified by having the 15 address bits that follow the I/G and U/L address bits set to zero. Certain values are used consistently throughout Token Ring LAN environments and, therefore, play a very similar role to standard group MAC addresses; these functional address values are also recorded in ISO/IEC TR 11802-2.

9.5 Bit-ordering and different MACs

NOTE—Throughout this subclause, considerations relating to the order of bit and/or octet transmission refer to the basic bit-serial model of transmission that applies to the representation of MAC frames at the boundary between the MAC sublayer and the Physical layer; see 6.2.3.

9.5.1 General considerations

The transmission of data on IEEE 802.3® and 802.4® LAN types is represented (6.2.3) as occurring LSB first within each octet. This is true for the entire frame: LAN MAC address fields (source and destination), MAC-specific fields (e.g., length field in IEEE 802.3® LANs), and the MAC Information field. (The MAC Information field is defined to be that part of the frame starting directly after the MAC header and ending immediately before the frame check sequence: e.g., LLC information, such as the Protocol Identification field, is contained in the MAC Information field.)

On some other LAN types, for which IEEE 802.5® is here used as the typical example, each octet of the MAC Information field is represented as being transmitted MSB first. The LAN MAC address fields (source and destination), however, are represented as being transmitted with the LSB of each octet first. Thus, the first bit transmitted is the I/G Address Bit, as on IEEE 802.3® and IEEE 802.4® LANs. For frames that originate within the MAC (e.g., MAC-embedded management frames), the ordering of bits within the MAC Information field is defined by the MAC specification—ISO/IEC 8802-5, etc.

For most purposes, the difference in the bit-orderings used to represent transmission of the octets of the MAC Information field is of no consequence, whether considered within a given MAC type, or across different MAC types. Each octet of user data is mapped to and from the appropriate ordering, symmetrically by the transmitting and receiving MAC entities. An unfortunate exception has occurred, however, where the octets concerned are those of a MAC address that is embedded, as user data, in the MAC Information field.

The consequences particularly affect the use of MAC addresses in mixed environments containing Token Ring LANs and non-Token-Ring LANs.

The following subclauses describe the problem and some of the issues arising from it.

9.5.2 Illustrative examples

This subclause illustrates the various bit- and octet-transmission scenarios that can occur, and it is intended as a basis for clarifying the issue of bit ordering for LAN MAC addresses across different MACs. Throughout, the examples make use of the OUI value AC-DE-48, introduced in 9.2.1. This three-octet value is considered in its two possible roles: as the first part of a five-octet protocol identifier, and as the first part of a six-octet LAN MAC address. The consistent representations of the OUI in its role as part of a protocol identifier are contrasted with the sometimes variable representations that apply to its role as part of a MAC address.

NOTE—Protocol identifiers always form part of the normal user data in a MAC Information field; hence, there is nothing special about OUI octets in their protocol identifier role.

For the examples, the bit significance of an OUI in general is defined to be as in Figure 10.

	MSB						LSB	
Octet 0	h	g	f	e	d	c	b	a
Octet 1	p	o	n	m	l	k	j	i
Octet 2	x	w	v	u	t	s	r	q

When used in LAN MAC addresses:
 Bit “a” of the OUI = I/G address bit. Bit “b” of the OUI = U/L address bit.

When used in protocol identifiers:
 Bit “a” of the OUI = M bit. Bit “b” of the OUI (always zero) = X bit.

Figure 10—Bit significance of an OUI

When transmitted on an IEEE 802.3® or IEEE 802.4® LAN (all data octets transmitted LSB first), the OUI portions of a protocol identifier and of a LAN MAC address appear as in Figure 11. When transmitted on an IEEE 802.5® LAN (data octets in the MAC Information field transmitted MSB first), the OUI portions of a protocol identifier, and of a LAN MAC address contained in a MAC Address field, appear as in Figure 12

In some circumstances, it is necessary to convey LAN MAC addresses as data within MAC Information fields, e.g., as part of a management protocol or a Network layer routing protocol.

For LAN types in which Figure 11 applies, such as IEEE 802.3®, the bit ordering within the octets of a LAN MAC address conveyed as data is the same as both the ordering when the address appears in a MAC address field and the ordering for octets of nonaddress information (see Figure 13).

For LAN types in which Figure 12 applies, such as IEEE 802.5® and FDDI, there appears to be a choice of representations for MAC addresses conveyed as data, as follows:

- The octets of the MAC address can be treated like any other data octets and transmitted with the bit-ordering of Figure 12 (A)
- The bit-ordering of Figure 12 (B) can be treated as a property of the MAC address rather than of the MAC address field as transmitted in MAC frames, and the MAC address octets can be transmitted with the bit-ordering reversed compared with normal data octets

(A) OUI within a protocol identifier (in MAC Information field, as normal data):



Figure 11—IEEE 802.3® , etc., frame: Order of bit and octet transmission for an OUI

(A) OUI within a protocol identifier (in MAC Information field, as normal data):

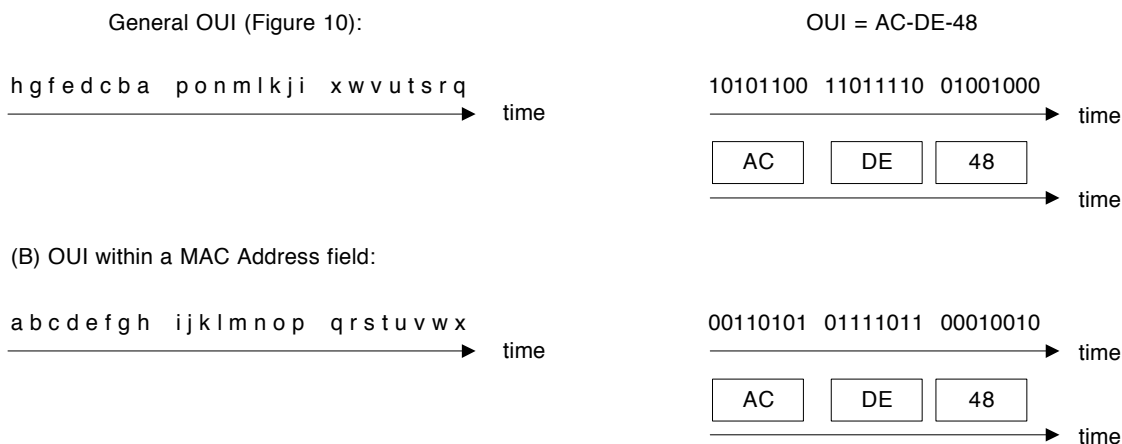


Figure 12—IEEE 802.5® , etc., frame: Order of bit and octet transmission for an OUI

OUI within a MAC address contained in the MAC Information field:



Figure 13—IEEE 802.3® , etc., frame: Order of bit and octet transmission for an OUI in a MAC address contained in a MAC Information field

In the case of IEEE 802.5® Token Ring LANs, approach b) was adopted early in the development and deployment of Token Ring technology. This has the unfortunate consequence that applications operating in environments containing a mixture of LAN types have to handle different representations of MAC addresses, according to the environment in which the MAC address is to be used; see Figure 14.

OUI within a bit-reversed MAC address contained in the MAC Information field:



Figure 14—IEEE 802.5®, etc., frame: Order of bit and octet transmission for an OUI in a MAC address contained in a MAC Information field (noncanonical format)

For other LAN types in which Figure 12 applies, including, in particular, FDDI, approach a) was adopted (Figure 15), at least in environments involving interconnection with IEEE 802.3®, and so on, LANs. However, where FDDI LANs are used in an IEEE 802.5® Token Ring environment, approach b) is used for consistency with the interconnected IEEE 802.5® LANs. In a mixed environment of FDDI, IEEE 802.3® and IEEE 802.5® LANs, frames constructed according to both approaches can occur on the FDDI LANs, at least; some care is needed in managing such an installation to avoid confusion between the formats.

OUI within a MAC address contained in the MAC Information field:




Figure 15—FDDI frame: Order of bit and octet transmission for an OUI in a MAC address contained in a MAC Information field (canonical format)

In Figure 11 through Figure 15, it can be seen that the interpretation of OUI bits as octet values is consistent in every case except Figure 14, in which the octet values are bit reversed. This reversal of the bit order applies only to all six octets (not just the OUI) of a MAC address placed in the MAC Information field of a frame by a protocol that uses the Token Ring Bit-reversed view of MAC addresses derived from Figure 12 (B). Frames containing, or possibly containing, such MAC addresses are described as having noncanonical format. Frames (on any LAN or LAN type) that cannot contain such MAC addresses are described as having canonical format.

Note that there is no way of knowing, from MAC layer information only, whether a particular frame is in canonical or noncanonical format. In general, this depends on which higher layer protocols are present in the frame.

9.5.3 Recommendation

Designers of protocols that operate above the Data Link layer are strongly recommended to avoid specifying new protocols that result in frames of noncanonical format, except where such a protocol is clearly an extension of existing practice in a strongly Token Ring environment.

10. Protocol discrimination above the MAC sublayer: Subnetwork Access Protocol (SNAP) and Ethernet types

10.1 Introduction

This clause outlines the mechanisms for the coexistence of multiple standard, public, and private network layer protocols within a single 802® station (10.2). It then describes the functions, features, and protocol format conventions for public and private protocols sharing a single LSAP (10.3). All public and private protocols using the IEEE 802® reserved LLC address assigned for public and private protocol use shall conform to this standard.

This clause further describes Ethernet types used to identify different protocols operating over the alternative Ethernet sublayer (10.4), and it describes the standard encapsulation specified for conveyance of such Ethernet-supported protocols on IEEE 802® LANs that do not intrinsically support an Ethernet sublayer (10.5).

A standard protocol is defined to be a protocol whose specification is published and known to the public but controlled by a standards body. A public protocol is defined to be a protocol whose specification is published and known to the public, but controlled by an organization other than a formal standards body. A private protocol is defined to be a protocol whose use and specification are controlled by a private organization.

By providing for the coexistence of multiple Network layer protocols, the migration of existing LANs to future standard protocols is facilitated, and multiple higher layer protocols are more easily accommodated.

10.2 Basic concepts

10.2.1 Coexistence of multiple protocols

Within a given layer, entities can exchange data by a mutually agreed upon protocol mechanism. A pair of entities that do not support a common protocol cannot communicate with each other. For multiple protocols to coexist within a layer, it is necessary to determine which protocol is to be invoked to process a service data unit delivered by the lower layer.

The following subclauses specify mechanisms for use when the LLC sublayer is present above the MAC layer, and when the alternative Ethernet sublayer is present above the MAC sublayer.

10.2.2 Multiple protocols above the LLC sublayer

Standard Network layer protocols have been assigned reserved LLC addresses, as recorded in ISO/IEC TR 11802-1. These addresses permit multiple standard network layer protocols to coexist at a single MAC station. One half of the LLC address space is reserved for such assignment.

Other protocols are accommodated in two ways. One way is by local assignment of LSAPs, for which the other half of the LLC address space is available. Thus, users can agree to use locally assigned LSAPs for either an instance of communication or a type of communication.

The second way is through the use of a particular reserved LLC address value that has been assigned for use in conjunction with the Subnetwork Access Protocol (SNAP, specified in 10.3), which provides for multiplexing and demultiplexing of public and private protocols among multiple users of a data link.

10.2.3 Multiple protocols above the Ethernet sublayer

The Ethernet MAC frame format includes a 16-bit type value, whose function is to identify the particular protocol pertaining to the user data contained in the frame. See 10.4 for further details.

10.3 Subnetwork Access Protocol

10.3.1 SNAP address

The reserved LLC address for use with SNAP is called the SNAP address. It is defined to be the bit pattern (starting with the LSB) Z1010101, in which the symbol Z indicates that either value 0 or 1 can occur, depending on the context in which the address appears (as specified in ISO/IEC 8802-2). The two possible values have Hexadecimal Representation AA and AB.

The SNAP address identifies, at each MSAP, a single LSAP for standard, public, and private protocol usage. To permit multiple public and private network layer protocols to coexist at one MSAP, each public or private protocol using SNAP must employ a protocol identifier that enables SNAP to discriminate among these protocols.

10.3.2 SNAP PDU format

Each SNAP PDU shall conform to the format shown in Figure 16 and shall form the entire content of the LLC information field.

In Figure 16, the Protocol Identification field is a five-octet field containing a protocol identifier whose format and administration are as described in 9.3. The Protocol Data field is a field whose length, format, and content are defined by a public or private protocol specification. Each public or private protocol begins its PDU format with the Protocol Identification field, which shall contain the protocol identifier assigned to the protocol.




Figure 16—SNAP PDU format

Figure 17 illustrates how a SNAP PDU appears in a complete MAC frame (the IEEE 802.3® MAC format is used for the example). The LLC control field (CTL) is shown for PDU type UI, Unnumbered Information, which is the most commonly used PDU type in this context; however, other information-carrying LLC PDU types may also be used with SNAP.




Figure 17—SNAP PDU in IEEE 802.3® MAC frame

10.4 Ethernet types: Format, function, and administration

The IEEE 802.3® MAC frame format is compatible with the alternative Ethernet MAC frame format, in the sense that frames of both formats can be freely intermixed on a given LAN and at a given LAN station. The service provided by use of the Ethernet frame format differs from the ISO/IEC 15802-1 MAC service in that there is a 16-bit type value associated with each frame transferred, and that the minimum amount of MAC-sublayer user data transferred is 64 octets.

An Ethernet type value is a sequence of two octets, interpreted as a 16-bit numeric value with the first octet containing the most significant 8 bits and the second octet containing the least significant 8 bits. Values in the range 0–1535 are not available for use.

The function of the Ethernet type value is to identify the protocol that is to be invoked to process the user data in the frame.

As with OUIs, administration of Ethernet types was originally undertaken by the Xerox Corporation, and it is now the responsibility of the IEEE using procedures defined by the IEEE RAC (see Clause 9). All assignments of Ethernet types made by the predecessor administration remain in effect under the IEEE's administration.

10.5 Encapsulation of Ethernet frames over LLC

This subclause specifies the standard method for conveying Ethernet frames across IEEE 802® LANs that offer only the LLC sublayer, and not the Ethernet sublayer, directly above the MAC sublayer.

An Ethernet frame conveyed on an LLC-only LAN shall be encapsulated in a SNAP PDU contained in an LLC PDU of type UI, as follows (see Figure 18):

- a) The Protocol Identification field of the SNAP PDU shall contain a protocol identifier in which
 - 1) The three OUI octets each take the value zero.
 - 2) The two remaining octets take the values, in the same order, of the two octets of the Ethernet frame's Ethernet type.
- b) The Protocol Data field of the SNAP PDU shall contain the user data octets, in order, of the Ethernet frame.
- c) The values of the Destination MAC Address and Source MAC Address fields of the Ethernet frame shall be used in the Destination MAC Address and Source MAC Address fields, respectively, of the MAC frame in which the SNAP PDU is conveyed.

NOTES

1—This encapsulation was originally specified in RFC 1042, which contains recommendations relating to its use. Further recommendations are contained in RFC 1390.

2—ISO/IEC TR 11802-5 (IEEE Std 802.1H-1997™) contains recommendations for bridges, addressing the consequences of certain problems that arose from differing interpretations of RFC 1042.

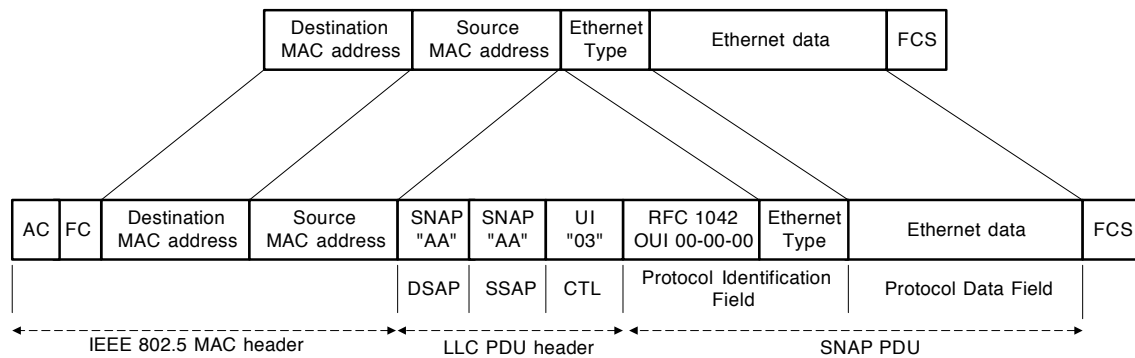


Figure 18—Ethernet frame SNAP-encapsulated in IEEE 802.5® frame

11. ISLAN and MAN support for isochronous bearer services

In addition to the mandatory LAN and MAN services described so far, some IEEE 802® LANs and MANs, notably, ISO/IEC 8802-9 and ISO/IEC 8802-6, also make provision for supporting isochronous bearer services. Isochronous bearer services are distinctive relative to packet services such as the MAC service and LLC, in that they maintain a flow of service data units at constant time intervals from a transmitter to a receiver for the duration of a service. In almost all circumstances, such isochronous bearer services are carried over duplex bidirectional connections thereby providing effective and efficient means of supporting the ubiquitous WAN voice telephony services. IEEE 802® ISLANs and MANs provide isochronous bearer services designed to interwork readily with these WAN services as standardized by ITU-T, in particular, as defined in the I series of ITU-T Recommendations.

NOTE—Use of an end-to-end physical-layer isochronous bearer service, which intrinsically delivers data with timing preserved, needs to be distinguished from use of a packet-based service for conveyance of time-sensitive data such as voice or video. The latter approach can be successful, given adequate bandwidth in the LANs and bridges, and provided that bridges do not introduce the possibility of excessive delay for packets carrying the time-sensitive data.

11.1 Key concepts

Applications requiring sustained periodic use of end-to-end network bandwidth are common. Two of the IEEE 802® standards address this requirement; both the ISO/IEC 8802-9 ISLAN standard and the optional isochronous service on an ISO/IEC 8802-6 Distributed Queue Dual Bus (DQDB) MAN use synchronously cyclic methods to ensure precise timing of the acceptance, transfer, and delivery of continuous streams of information. This applies whether or not asynchronous packet services are also provided concurrently.

The ISLAN employs a time-division multiplex (TDM) bearer mechanism within the Physical layer. The isochronous service on a DQDB MAN uses a Pre-Arbitrated function to ensure precisely timed access to the media, as distinct from the packet-service Queued Arbitrated function. In each case, this permits the support and delivery of a plurality of transparent isochronous service channels, the provision of an octet alignment signal for these channels, and a facility to provide and accept these precise timing signals. It is the provision of the timing signal that principally distinguishes the isochronous services from the asynchronous packet services that are provided from the ISLAN Physical layer.

Both methods for providing isochronous bearer services require the prior establishment of a connection for each isochronous information flow. For the DQDB MAN isochronous services, the mechanisms used for establishing and clearing connections are left outside of the scope of the IEEE 802® Standard. For ISLAN, see 11.4 below for an overview of the mechanisms used.

11.2 Applications

Applications for isochronous bearer services include the following:

- PBX interconnections at DS1 (1.544 Mbit/s) or E1 (2.048 Mbit/s) rates
- Low (384 kbit/s) to medium (44.2097 Mbit/s) bandwidth constant bit-rate compressed video
- Channels with bandwidth in multiples of 56 or 64 kbit/s for conveyance of voice and audio
- Multimedia combinations of these along with data services

Multimedia applications can require simultaneous, integrated use of two or more of these audio/voice, video, text, and graphics information streams. This can require the conveyance over common bearer channels of multiple isochronous and bursty information streams, in distinct channels with specific timing relationships. This provides the main rationale for incorporating isochronous services in LANs and MANs.

The provision of isochronous (TDM bearer) services enables direct interworking with a WAN network such that the access unit or hub (LAN/WAN interworking unit) can forward information in its native form, so that ISLAN terminals and access units do not have to provide gateway functions to transform the user information streams. In addition, the provision of isochronous bearer services greatly reduces the latency as information is queued/dequeued at the Physical service interface. This is of value in meeting established norms of loop delay on end-to-end connections for many human interactive services.

Typical isochronous services include the following:

- Unrestricted 64 kbit/s information transfer
- Restricted 56 kbit/s information transfer
- Synchronous data
- Facsimile data
- Wideband video and image transfer

The ISO/IEC 8802-9 ISLAN is specifically designed to provide concurrent support for LLC conformant packet services and narrowband ITU-T conformant ISDN services, both packet and isochronous, as defined in ITU-T I-series Recommendations.

11.3 Isochronous service access points and PhSAPs

The ISLAN and DQDB MAN standards both support concurrent packet and isochronous services within the Physical layer by means of convergence functions. In both cases, there is therefore a need for explicit identification of the distinct packet and isochronous channels provided over the common media supporting the ISLAN or MAN. In addition, both offer support for multiple isochronous bearer channels, and thus, there is a need to distinguish among multiple Physical service connections.

In the ISLAN standard, multiple PhSAPs are used to provide for the plurality of connections at the Physical service boundary, and for access by the user to the Physical-layer services for packet transfer and for D channel signaling (see Figure 19). For the DQDB MAN isochronous services, Connection Endpoint (CEs) identifiers are used for a similar purpose.




Figure 19—ISLAN RM (end station)

PhSAPs and CEs are the architectural mechanism by which symbol streams are passed to the Physical service provider by the Physical service user, and to the Physical service user by the Physical service provider. The distinction of different PhSAPs at a single ISLAN Physical service boundary is required because that service boundary is used simultaneously to provide both packet-mode and (multiple) circuit-mode Physical services.

It is a function of the layer management of the PHY multiplexing sublayer to provide each Physical service user with both the information stream and a channel identifier that is mapped onto the PhSAP relevant to the service provided to that user.

11.4 ISLAN signaling

The ISLAN standard provides for direct interworking with ITU-T conformant ISDN I-series Recommendations. These require means of establishing, maintaining, and disestablishing end-to-end connections across (in the ISLAN case) both the ISLAN and intervening WAN. To this end, ISO/IEC 8802-9 includes specifications for extensions to the ITU-T I- and Q-series signaling protocols carried in a signaling-specific D channel. This is a packet-based protocol, but to achieve interworking with the ITU-T Recommendations, it is distinct from the LLC packet service provided over ISLAN.

The signaling service provided by the protocols carried in the D channel permits negotiation of end-to-end network resources such that a guaranteed service can be provided to the users of an end-to-end isochronous channel. Thus, the ISLAN management signalling entity is responsible, as a management agent, for performing the following tasks:

- Negotiation of bandwidth (configuration) management, fault management, performance management, and security management of the access link between device and hub
- Negotiation of service provisioning over the local ISLAN interface in order to access the WAN-based ISDN services

The ISLAN signaling message elements are provided in compliance with international (ITU-T) network signaling methods with appropriate protocol extensions.

Annex A

(informative)

Bibliography (Additional references for LAN/MAN-related standards)

A.1 General LAN/MAN standards and specifications

[B1] *Ethernet Version 2.0, A Local Area Network—Data Link Layer and Physical Layer Specifications*. Digital Equipment Corp., Intel Corp., and Xerox Corp., November 1982.

[B2] IEEE 100™, The Authoritative Dictionary of IEEE Standards Terms, Seventh Edition.

[B3] IEEE Std 802.10-1992™, IEEE Standards for Local and Metropolitan Area Networks: Interoperable LAN/MAN Security (SILS) (includes IEEE Std 802.10b-1992™).⁶

[B4] ISO/IEC TR 8802-1:1997, Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Technical reports and guidelines—Part 1: Overview of Local Area Network Standards.

[B5] ISO/IEC 8802-3:1999 (IEEE Std 802.3-1998™), Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications.

[B6] ISO/IEC 8802-4:1990 (IEEE Std 802.4-1990™), Information processing systems—Local area networks—Part 4: Token-passing bus access method and physical layer specifications.

[B7] ISO/IEC 8802-5:1998 (IEEE Std 802.5-1998™), Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 5: Token ring access method and physical layer specifications.

[B8] ISO/IEC 8802-6:1994 (IEEE Std 802.6-1994™), Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 6: Distributed Queue Dual Bus (DQDB) access method and physical layer specifications.

[B9] ISO/IEC 8802-9:1996 (IEEE Std 802.9-1996™), Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 9: Integrated Services (IS) LAN Interface at the Medium Access Control (MAC) and Physical (PHY) Layers.

[B10] ISO/IEC 8802-11:1999 (IEEE Std 802.11-1999™), Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications.

⁶IEEE publications are available from the Institute of Electrical and Electronics Engineers, 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331, USA (<http://standards.ieee.org/>).

[B11] ISO/IEC 8802-12:1998 (IEEE Std 802.12-1998[™]), Information technology — Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 12: Demand-Priority access method, physical layer and repeater specifications.

[B12] ISO 9314-1:1989, Information processing systems—Fibre Distributed Data Interface (FDDI)—Part 1: Token Ring Physical Layer Protocol (PHY).

[B13] ISO 9314-2:1989, Information processing systems—Fibre Distributed Data Interface (FDDI)—Part 2: Token Ring Media Access Control (MAC).

[B14] ISO 9314-3:1990, Information processing systems—Fibre Distributed Data Interface (FDDI)—Part 3: Token Ring Physical Layer Medium Dependent (PMD).

[B15] ISO 9314-6:1989, Information processing systems—Fibre Distributed Data Interface (FDDI)—Part 6: Token Ring Station Management (SMT).

[B16] ISO/IEC TR 11802-5:1997 (IEEE Std 802.1H-1997[™]), Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Technical reports and guidelines—Part 5: Media Access Control (MAC) Bridging of Ethernet V2.0 in Local Area Networks.

[B17] ISO/IEC 15802-5:1998 (IEEE Std 802.1G-1998[™]), Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Common specifications—Part 5: Remote Media Access Control (MAC) Bridging.

[B18] RFC 1042, A Standard for the Transmission of IP Datagrams over IEEE 802[®] Networks. Postel, J., and Reynolds, J., February 1988.⁷

[B19] RFC 1390, Transmission of IP and ARP over FDDI Networks. Katz, D., January 1993.

A.2 Management

NOTE—Many of the standards listed in Clause 2 and A.1 contain specifications of their managed objects.

[B20] IEEE Std 802.1F-1993[™], IEEE Standards for Local and Metropolitan Area Networks: Common Definitions and Procedures for IEEE 802[®] Management Information.

NOTE—Some of the content of IEEE Std 802.1F-1993[™] is also included in ISO/IEC 10742.

[B21] ISO/IEC 7498-4:1989, Information processing systems—Open Systems Interconnection—Basic Reference Model—Part 4: Management framework.

[B22] ISO/IEC 8824-1:1995, Information technology—Abstract Syntax Notation One (ASN.1)—Specification of basic notation.

[B23] ISO/IEC 8824-2:1995, Information technology—Abstract Syntax Notation One (ASN.1)—Information object specification.

⁷Internet RFCs are available from the Internet Engineering Task Force on the World Wide Web browser at <http://www.ietf.org/rfc/rfcnnnn.txt> (where nnnn is the four-digit RFC number, padded with leading zeros if necessary). For more information, call the IETF secretariat at +1-703-620-8990.

[B24] ISO/IEC 8824-3:1995, Information technology—Abstract Syntax Notation One (ASN.1)—Constraint specification.

[B25] ISO/IEC 8824-4:1995, Information technology—Abstract Syntax Notation One (ASN.1)—Parameterization of ASN.1 specifications.

[B26] ISO/IEC 8825-1:1995, Information technology—ASN.1 encoding rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER).

[B27] ISO/IEC 9595:1991, Information technology—Open Systems Interconnection—Common management information service definition.

[B28] ISO/IEC 9596-1:1991, Information technology—Open Systems Interconnection—Common management information protocol—Part 1: Specification

[B29] ISO/IEC 10040:1992, Information technology—Open Systems Interconnection—Systems management overview.

[B30] ISO/IEC 10164-1/8:1993, Information technology—Open Systems Interconnection—Systems management: [various functions].

[B31] ISO/IEC 10165-1:1993, Information technology—Open Systems Interconnection—Structure of management information: Management information model.

[B32] ISO/IEC 10165-2:1992, Information technology—Open Systems Interconnection—Structure of management information: Definition of management information.

[B33] ISO/IEC 10165-4:1992, Information technology—Open Systems Interconnection—Structure of management information: Guidelines for the definition of managed objects.

[B34] ISO/IEC 10742:1994, Information technology—Telecommunications and information exchange between systems—Elements of management information related to OSI Data Link Layer standards.

[B35] ISO/IEC 15802-2:1995 (IEEE Std 802.1B-1995™), Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Common specifications—Part 2: LAN/MAN management, service and protocol.

[B36] ISO/IEC 15802-4:1994 (IEEE Std 802.1E-1994™), Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Common specifications—Part 4: System load protocol.

[B37] RFC 1095, The Common Management Services and Protocol over TCP/IP (CMOT). Warrior, U., and Besaw, L., April 1989.

[B38] RFC 1155, Structure and Identification of Management Information for TCP/IP-based Internets. Rose, M., and McCloghrie, K., May 1990.

[B39] RFC 1157, A Simple Network Management Protocol (SNMP). Case, J., Fedor, M., Schoffstall, M., and Davin, J., May 1990.

[B40] RFC 1212, Concise MIB Definitions. Rose, M., and McCloghrie, K. (Editors), March 1991.

[B41] RFC 1215, A Convention for Defining Traps for use with the SNMP. Rose, M., March 1991.

IEEE
Std 802-2001[®]

[B42] RFC 1493, Definitions of Managed Objects for Bridges. Decker, E., Langille, P., Rijsinghani, A., and McCloghrie, K., July 1993.

[B43] RFC 1643, Definitions of Managed Objects for the Ethernet-like Interface Types. Kastenholz, F., July 1994.

[B44] RFC 1902, Structure of Management information for Version 2 of the Simple Network Management Protocol (SNMPv2). Case, J., McCloghrie, K., Rose, M., and Waldbusser, S., January 1996.

[B45] RFC 1903, Textual Conventions for Version 2 of the Simple Network Management Protocol (SNMPv2). Case, J., McCloghrie, K., Rose, M., and Waldbusser, S., January 1996.

[B46] RFC 1904, Conformance Statements for Version 2 of the Simple Network Management Protocol (SNMPv2). Case, J., McCloghrie, K., Rose, M., and Waldbusser, S., January 1996.

[B47] RFC 1905, Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2). Case, J., McCloghrie, K., Rose, M., and Waldbusser, S., January 1996.